834-E
NS 1.10/7:985

DIVISION OF COMPUTER RESEARCH

SUMMARY OF AWARDS

Fiscal Year 1985

COMPLETED
109

NATIONAL SCIENCE FOUNDATION
Directorate for Mathematical and Physical Sciences
Washington, DC 20550

BEST COPY AVAILABLE

The Foundation provides awards for research in the sciences and engineering. The awardee is wholly responsible for the conduct of such research and preparation of the results for publication. The Foundation, therefore, does not assume responsibility for such findings or their interpretation.

The Foundation welcomes proposals on behalf of all qualified scientists and engineers, and strongly encourages women and minorities to compete fully in any of the research and research-related Programs described in this document.

In accordance with Federal statutes and regulations and NSF policies, no person on grounds of race, color, age, sex, national origin, or physical handicap shall be excluded from participation in, denied the benefits of, or be subject to discrimination under any program or activity receiving financial assistance from the National Science Foundation.

NSF has TDD (Telephonic Device for the Deaf) capability which enables individuals with hearing impairments to communicate with the Division of Personnel Management for information relating to NSF programs, employment, or general information. This number is (202) 357-7492.

BEST

DIVISION OF COMPUTER RESEARCH

SUMMARY OF AWARDS

Fiscal Year 1985

NATIONAL SCIENCE FOUNDATION
Directorate for Mathematical and Physical Sciences
Washington, DC 20550

BLANK PAGE

Preface

The purpose of this report is to provide the scientific community with a summary of those grants awarded by the Division of Computer Research in Fiscal Year 1985. The report lists fewer than the actual number of projects currently receiving support because the duration of some grants exceeds one year.

Similar areas of research are grouped together for reader convenience. The reader is cautioned, however, not to assume these categories represent either the totality of interests of each program or the total scope of each grant. Projects may bridge several programs or deal with topics not explicitly mentioned herein. Thus, these categories have been assigned administratively and for the purposes of this report only.*

In this document, grantee institutions and principal investigators are identified first. Award identification numbers, award amounts, and award durations are enumerated after the individual project titles. Within each category, awards are listed alphabetically by state and institution.

Readers wishing further information on any particular project described in this report are advised to contact principal investigators directly.

Kent K. Curtis
Division Director
Division of Computer Research

*To get an idea of the scope of possible programs, consult *What Can Be Automated? The Computer Science and Engineering Research Study (COSERS)*, Bruce Arden (Ed.), MIT Press, Cambridge, MA (1980).

Table of Contents

	Page
Preface	iii
Table of Contents	iv
Introduction	vi
Staff	vii
Advisory Committee	viii
Summary	ix
Index of Principal Investigators	93
Computational Mathematics	1
Numerical Methods	1
Algorithms and Software	4
Symbolic and Algebraic Manipulation	7
Computer Systems Design	9
VLSI Design Methodology	9
Computer System Architecture	10
Fault Tolerance and Reliability	13
Computer System Performance Measurement and Evaluation	14
Computer Graphics	14
Intelligent Systems	17
Computer Vision and Image Processing	17
Natural Language and Signal Understanding	21
Automatic Theorem Proving, Concept Learning, and Inference	23
Knowledge Representation and Problem Solving	27
Software Engineering	33
Software Studies	33
Program Testing and Fault Tolerance	34
Software Tools and Programming Environments	35
Compilers and Language Implementations	37
Software Systems Science	39
Programming Languages: Design, Semantics, Implementation	39
Operating Systems and Concurrency	42
Data Bases	42
Programming Methodology	44
Distributed Computing	47
Verification	49
Special Projects	51
Privacy & Security	51
The Societal Impact of Computing	52
Modelling & Simulation	52
Networking, Computer Communications, & Distributed Computing	53
Other Projects, Symposia, Colloquia, and Studies	56
Theoretical Computer Science	59
Concrete Complexity and Analysis of Algorithms	59
Parallel, Distributed, and Data Base Models	66

Foundations of Computer Science 70
Automata and Language-based Models 73

Computer Research Equipment 75

Coordinated Experimental Research 77
Experimental Computer Research 77
Computer Science Research Network (CSNET) 83

Presidential Young Investigator Awards 85

Small Business Innovation Research 91

BEST C...

Introduction

This report provides summaries of awards made by the National Science Foundation (NSF) through the Division of Computer Research in Fiscal Year 1985. In Fiscal Year 1985, the Division of Computer Research award activities were carried out through eight programs. Topics of interest to the Foundation include, but are not limited, to the following:

Computer Systems Design

Principles of computer systems design relating to the structure of computer systems or the process of systems design. Topics include, but are not limited to: computer system architecture; distributed computer systems; integrated hardware software systems; performance measurement and evaluation; fault tolerant systems; logic design; computer graphics; man-machine interaction; and VLSI design methodology. The scope of this program includes experimental implementation where that is an integral part of the research.

Intelligent Systems

Computer-based systems that have some of the characteristics of intelligence. Relevant areas of inquiry include: knowledge engineering; automated theorem-proving; mechanical inferencing; problem solving; pattern analysis; computer vision; natural language and speech understanding; and areas related to the automatic analysis and handling of complex tasks.

Software Engineering

The structure and design process of computer software, especially verification, testing, portability, reliability, and human interfacing to numeric and non-numeric software systems. Areas of emphasis include: program validation and testing; software tools; and human factors in software design and use. The program also supports research in computationally oriented numerical analysis, the design and construction of high quality portable software for scientific research, and experimental implementation where that is an integral part of the research.

Software Systems Science

Conceptual basis for the specification of future software systems and the necessary experimentation with such systems, including: advanced programming languages and optimizing compilers; functional and relational specification; program transforming systems; systems for the verification and proof of correctness of programs; the study of the concurrency of operations; the discovery of new algorithms and improved measures of effectiveness of known algorithms.

Special Projects

New directions in computer science and applications including computer networks, database management, computer-based modeling, privacy and security of computer systems, and other topics of special interest in computer science.

Theoretical Computer Science

Theories of computation and formal languages, analysis of algorithms; theoretical models for computation, and other theoretical problems concerned with the foundations of computer science.

Computer Research Equipment

Support of special purpose equipment needed by more than one computer research project and difficult to justify on a single project.

Coordinated Experimental Research

Experimental Computer Research — The establishment and enhancement of experimental research facilities, needed technical and professional support personnel, and necessary maintenance of facilities; large multi-investigator projects in experimental computer research (5 year awards expected).

Computer Science Research Network (CSNET) — Build a logical computer-based communication network spanning ARPANET, public packet networks, and PHONENET, a telephone-based relay system; provide communication services, file transfer, and access to remote databases; stimulate new classes of network research activities.

Additional Information

For additional information on any of the projects, contact the principal investigators directly.

Division of Computer Research

Kent K. Curtis
Division Director
(202) 357-9747

W. Richards Adrion
Deputy Division Director
(202) 357-1185

Program or Area	Program Director	Telephone
Computational Mathematics	Bruce H. Barnes	(202) 357-7345
Computer Research Equipment	John R. Lehmann	(202) 357-7349
Computer Systems Design	John R. Lehmann	(202) 357-7349
Coordinated Experimental Research	Harry G. Hedges	(202) 357-7349
	Robert C. Minnick	(202) 357-7349
CSNET	Harry G. Hedges	(202) 357-7349
Intelligent Systems	Y. T. Chien	(202) 357-7345
Software Engineering	Bruce H. Barnes	(202) 357-7345
Software Systems Science	Thomas A. Keenan	(202) 357-7375
Special Projects	Lawrence H. Oliver	(202) 357-7375
Theoretical Computer Science	Thomas J. Head	(202) 357-7375

CSNET ADDRESS: last-name @ nsf

November 1985

Division of Computer Research

Advisory Committee
November 1985

Dr. Clarence A. Ellis
Microelectronics & Computer Technology Corporation
(512) 834-3384
ARPANET: ellis@mcc.arpa

Dr. Charles W. Gear
University of Illinois
(217) 333-0195
ARPANET: gear@.cs.uiuc.edu

Dr. A. Nico Habermann
Carnegie-Mellon University
(412) 578-2592
ARPANET: habermann@c.cs.cmu.edu

Dr. John E. Hopcroft
Cornell University
(607) 256-4052
CSNET: hopcroft@cornell

Dr. Kenneth W. Kennedy, Chairman
Rice University
(713) 527-4834
CSNET: ken@rice
(1985-86) Stanford University
(415) 497-0474

Dr. Rao Kosaraju
The John Hopkins University
(301) 338-7018
CSNET: srk@jhu

Dr. Lawrence H. Landweber
University of Wisconsin
(608) 262-1204
ARPANET: lhl@rsch.wisc.edu

Dr. Wendy G. Lehnert
University of Massachusetts
(413) 545-3639
CSNET: lehnert@umass

Dr. Nancy A. Lynch
Massachusetts Institute of Technology
(617) 253-7225
ARPANET: lynch@mit-xx.arpa

Dr. John T. Pinkston, III
Microelectronics & Computer Technology Corporation
(512) 834-3544
ARPANET: pinkston@mcc.arpa

Dr. Burton J. Smith
Supercomputing Research Center, IDA
(301) 731-8065
No NET Address

Dr. Lawrence Snyder
University of Washington
(206) 543-9265
ARPANET: snyder@washington.arpa

Summary

	Number of Projects	Value of Awards
Computational Mathematics	37	$1,433,290
Numerical Methods	18	$678,669
Algorithms and Software	14	$506,479
Symbolic and Algebraic Manipulation	5	$248,142
Computer Systems Design	30	$3,109,196
VLSI Design Methodology	5	$406,308
Computer System Architecture	13	$1,355,980
Fault Tolerance and Reliability	3	$363,796
Computer System Performance Measurement and Evaluation	4	$175,048
Computer Graphics	5	$808,064
Intelligent Systems	66	$5,023,107
Computer Vision and Image Processing	19	$961,747
Natural Language and Signal Understanding	8	$505,283
Automatic Theorem Proving, Concept Learning, and Inference	16	$1,267,955
Knowledge Representation and Problem Solving	23	$2,288,122
Software Engineering	29	$2,041,965
Software Studies	7	$340,518
Program Testing and Fault Tolerance	4	$277,270
Software Tools and Programming Environments	11	$917,340
Compilers and Language Implementations	7	$506,837
Software Systems Science	45	$3,933,685
Programming Languages: Design, Semantics, Implementation	14	$1,178,211
Operating Systems and Concurrency	3	$154,932
Data Bases	8	$771,819
Programming Methodology	12	$995,841
Distributed Computing	7	$780,303
Verification	1	$52,579
Special Projects	32	$994,244
Privacy & Security	4	$147,285
The Societal Impact of Computing	1	$54,855
Modelling & Simulation	2	$137,162
Networking, Computer Communications, & Distributed Computing	15	$407,532
Other Projects, Symposia, Colloquia, and Studies	10	$247,410
Theoretical Computer Science	86	$4,330,764
Concrete Complexity and Analysis of Algorithms	40	$2,030,174
Parallel, Distributed, and Data Base Models	21	$1,297,771
Foundations of Computer Science	19	$659,518
Automata and Language-based Models	6	$343,298
Computer Research Equipment	15	$1,464,248

	Number of Projects	Value of Awards
Coordinated Experimental Research	25	$14,991,601
Experimental Computer Research	24	$14,762,900
Computer Science Research Network (CSNET)	1	$228,701
Presidential Young Investigator Awards	25	$1,265,817
Small Business Innovation Research	5	-0-

Note. The above amounts are those supported by the Division of Computer Research and do not reflect the total of the awards. Refer to the individual summaries for details.

Computational Mathematics

Numerical Methods

California Institute of Technology. Heinz-Otto Kreiss. *Numerical Methods for Nonlinear Ordinary Differential Equations* (*Mathematical Sciences and Computer Research*) (DMS-8312264 A01), $81,700, 12 mos. (Joint support with the Applied Mathematics Program and the Office of Advanced Scientific Computing - Total Grant $110,400)

The research focuses on analytical studies and numerical experiments to further develop robust and efficient procedures. The class of systems of ODE's are those that are classified essentially non-oscillatory stiff ODE's and high oscillatory stiff ODE's. A prototype computer code is being further developed for application to realistic problems such as reentry-roll resonance.

Stanford University. Gene H. Golub. *Numerical Analysis and Scientific Computing* (*Computer Research*) (DCR-8412314), $170,884, 12 mos.

This numerical analysis research project centers in the areas of linear algebra and ordinary differential equations. The linear algebra problems are often motivated by problems which arise from differential equations. In particular, Wilkinson is studying the sensitivity of eigen-systems with respect to perturbation in the matrix of coefficients. The recent work of Golub has been motivated by the desire to solve elliptic equations over irregular domains. Here again it is necessary to study this problem from both the linear algebra and differential equations viewpoints. Dahquist has recently devised a new procedure for deflating companion matrices and from that there are expectations concerning contractivity analysis.

Aside from the theoretical developments in numerical analysis, the results should lead to new algorithms and software. Even with the advent of super-computers, there is still a need for the development of new algorithms. No matter how fast the super-computer, there is always another problem which requires the development of new algorithms and their analysis.

University of California - San Diego. James R. Bunch. *Stability of Matrix Computations for Toeplitz Systems* (*Mathematical Sciences and Computer Research*) (DMS-8318412A01), $16,325, 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $32,650)

In this project Professor Bunch is studying the stability of matrix computation for Toeplitz matrices, that is, matrices where the elements on each diagonal are all equal. The solution of such matrices arise out of problems involving time series analysis, linear filtering, image processing, control theory, partial differential equations and approximation theory.

Professor Bunch is studying algorithms for n x n matrices that attempt to improve on the classical work estimate of n^3 operations by using a QR decomposition.

University of Colorado - Boulder. Richard H. Byrd. *Trust Region Methods for Minimization* (*Computer Research*) (DCR-8403483 A01), $48,432, 12 mos.

This research is concerned with constrained and unconstrained optimization, solving systems of nonlinear equations, and nonlinear least squares. Four topics are being investigated. First, the development of tensor methods for nonlinear equations and their extention to unconstrained optimization is continued. These methods show great promise of yielding general purpose algorithms that are highly efficient on singular and nonsingular problems. Second, practical local methods for constrained optimization that are one-step q-superlinearly convergent under realistic assumptions are being developed. Third, global methods for constrained optimization that are practical and satisfy global convergence theory and that perform robustly even from poor starting values are being developed. Finally, algorithms for several important generalizations of nonlinear least squares not addressed by currently available software are being investigated. These include data fitting problems where there are errors in the independent variables, modeling problems where the model is implicit, and curve fitting problems where an orthogonal measure of the distance from the data points to the fitting function is desired.

Purdue University. Walter Gautschi. *Applied Orthogonal Polynomials* (*Computer Research and Mathematical Sciences*) (DCR-8320561 A01), $35,183, 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $50,263)

The principal objective of this research is a study of constructive methods for Gauss-type quadrature rules and, more generally, for orthogonal polynomials, assuming arbitrary weight distributions. An additional goal is the preparation of related mathematical software. The study of Gauss-type quadrature rules in-

cludes Radau- and Lobatto-type formulae, as well as Gauss formulae with multiple nodes. In addition to developing methods for generating general orthogonal polynomials, specific problems connected with modifying the weight distribution are being investigated.

University of Maryland: Ivo Babuska, *Higher Order Finite Element Methods and Adaptive Approaches (Mathematical Sciences and Computer Research)*, (DMS-8315316 A01), $26,470, 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $79,500)

In this project Professor Babuska is investigating adaptive order finite element methods (p and h-p versions) for partial differential equations arising in solid mechanics. This research is being conducted in close collaboration with investigators from the Center of Computational Mechanics, Washington University, St. Louis, Missouri.

Professor Babuska is providing guidance on numerical experiments using a software package, FIESTA (Software System for Static Analysis of Solid Structures Based on the p-Version of the Finite Element Method).

University of Minnesota: J. B. Rosen: *Global Optimization For Large Scale Problems Using Vector Processing*, (DCR-8405489 A01 &A02), $27,167 + 10 hours of supercomputer time, 12 mos. (Joint support with the Office of Advanced Scientific Computing - Total Grant $53,167).

The objective of this research is to develop and implement efficient and robust algorithms for the computational solution of large scale global optimization problems. The problems considered include the minimization of a concave function subject to linear constraints, the minimization of an indefinite quadratic function subject to linear constraints, and minimum cost network flow problems with some concave costs. All these problems are combinatorial in nature. Large scale problems of this kind typically contain many variables which appear only linearly. In order to solve such large scale problems the algorithms must be designed to take advantage of the linearity by restricting the combinatorial aspects to a subspace of nonlinear variables.

There are many important application areas which can be formulated as problems of the type considered here. These include a wide class of scheduling, investment and planning models. More recently, important aspects of chip design and physical database design problems have been formulated in these terms.

The algorithms are being developed, implemented and tested, using a Class VI vector processing machine, so that thorough testing can be done in a reasonable amount of time. The algorithms and software being developed take advantage of the vector processing and parallel capabilities of such machines, so that they can be used efficiently on the new generation of supercomputers.

University of New Mexico: Robert Russell, *Numerical Methods for Boundary Value Problems in ODEs, with Applications to PDE's (Mathematical Sciences and Computer Research)*, (DMS-8401019 A01), $16,200, 12 mos. (Joint support with the Applied Mathematics Program - Total Grant $32,400)

Professor Russell is studying a variety of numerical methods for boundary value problems for ordinary differential equations (BVODE's). The underlying approach is to carefully analyze the associated discretization matrix factorizations. In this framework, he is deriving improved stability results for invariant imbedding, study conditions of collocation matrices using spline bases, and re-interpret collocation and initial value methods for singular perturbations.

Professor Russell is a co-author with U. Ascher and J. Christiansen of COLSYS (a collecation of software for BVODE's computer program) - a popular software code. In this project, Professor Russell is modifying COLSYS to incorporate a more efficient linear system solver and a more robust spline basis representation. The BVODE techniques are being used to study initial value and boundary value partial differential equations.

New York University: James Demmel, *Stably Computing the Kronecker Canonical Form for Uncertain Data (Computer Research and Mathematical Sciences)*, (DCR -8501708), $20,650, 24mos. (Joint support with the Applied Mathematics Program - Total Grant $41,300).

The generalized eigenvalue problem $Ax = zBx$, where A and B are given matrices, x is the eigenvector to be computed and z is the eigenvalue to be computed, arises in problems of systems theory, structural analysis, and numerical analysis. In applications unavoidable uncertainty in the entries of A and B causes uncertainty in x and z. Previous analyses of the uncertainties in x and z are unsatisfactory for two reasons. First, they deal only with very small perturbations in A and B. Second, they do not deal with the singular case, i.e., when $Ax = zBx$ has a solution vector x for every z. This pathological situation is of particular interest in systems theory, when analyzing the controllability and observability of a control system. Programs previously developed for computing x and z reflect these weaknesses in the theory by producing results without sufficient indication of their reliability. This research focuses on extending analyses of uncertainty to cover the cases of large perturbations and singular pencils. The approach involves partitioning the eigenvalues of $Ax = zBx$ into groups whose deflating subspaces are insensitive to perturbations in A and B even though the individual eigenvalues and eigenvectors may be very sensitive. This approach has been developed so far

only for the nonsingular case, the singular case being more unstable and difficult.

New York University; Michael Overton; *Numerical Methods for Optimization (Computer Research and Mathematical Science)*; (DCR-8502014); $22,000; 12 mos; (Joint support with the Applied Mathematics Program - Total Grant $42,000).

This research is investigating numerical methods for solving optimization problems, addressing itself to four rather disjointed topics: (1) Analysis of nonlinear programming methods which update an approximation to a projected Hessian matrix, (2 and 3) relate to minimizing a sum of Euclidean norms, and (4) numerical methods for the inverse eigenvalue problem. In all cases the goals are an improved mathematical framework, efficient and reliable numerical methods, and general purpose, user-oriented software.

New York University; Olof Widlund; *Fast Iterative Methods for Certain Finite Element and Queuing Network Problems*; (DCR-8405506 A01); $31,750; 12 mos; (Joint support with the Applied Mathematics Program - Total Grant $63,500).

This project involves the development of iterative methods for the numerical solution of several types of special large linear algebra problems. Work continues on iterative substructuring (domain decomposition) techniques for elliptic finite element problems. The main issue is now the extension of methods which have proven highly successful for second order problems in the plane to other elliptic systems. Related component mode substitution methods, which have been developed by structural engineers as an alternative to traditional finite element eigenvalue methods, are being investigated. Another main part of the project is the application of similar numerical techniques to certain queuing network problems. In these problems the Kolmogorov balance equations are generated and the steady state probability distribution is computed in terms of the null vector of an often very large, sparse nonsymmetric matrix. This matrix has much in common with those arising in finite difference problems and differs from decomposable problems by matrix of relatively low rank. A systematic study of the use of various iterative methods and special fast solvers and of the underlying continuous problems is being undertaken.

North Carolina State University - Raleigh; Robert Plemmons and Carl Meyer; *Numerical Linear Algebra (Mathematical Sciences and Computer Research)*; (DMS-8219500 A03); $10,300; 12 mos; (Joint support with the Applied Mathematics Program and the Office of Advanced Scientific Computing - Total Grant $55,900.

Professors Plemmons and Meyer are applied mathematicians working in the area of numerical linear algebra. They are working on a variety of problem that arise in the study of M-matrices - a class of matrices that arise in contexts of solving partial differential equations and queueing theory.

University of Pittsburgh; Werner C. Rheinboldt; *Computational Methods for Parametrized Equations (Computer Research)*; (DCR-8309926 A01); $59,669; 12 mos.

This research concerns the development and study of computational methods for the analysis of the solution manifolds of parametrized equations. Such problems arise frequently in applications, for instance, in structural mechanics, fluid flow, network analysis, etc. Some of the topics under consideration concern the continued development and further extension of a general program package for the trace of paths on given solution manifolds, an investigation from the viewpoint of numerical analysis of the reduced basis technique pioneered by structural engineers, a study of a new approach combining continuation and discrete variable methods for the solution of differential equations on manifolds, and further work on the development and application of estimates of the discretization error of parametrized equations.

Rice University; Robert E. Bixby; *Efficient Detection and Solution of Partial-Network Linear Programs (Computer Research)*; (DCR-8416187); $36,000; 12 mos.

This research concerns linear programs (LPs) that have a network structure hidden within some subsection of their constraints. Such LPs can be transformed to partial-network LP's having activities with simple interpretations in terms of flows and transshipments, and solutions that can be obtained efficiently by the use of specialized simplex methods. This research emphasizes:

1. Algorithms that find embedded partial networks by scaling the rows and columns of the LP constraint matrix,
2. Algorithms that find hidden partial networks through column scaling and elementary row operations on the LP constraint, and
3. Specialized simplex methods that take advantage of partial networks.

University of Wisconsin - Madison; Olvi Mangasarian; *Computational Optimization and Large Scale Systems (Computer Research)*; (DCR-8420963); $24,510; 12 mos; (Joint support with the Air Force Office of Scientific Research - Total Grant $74,510).

Professor Mangasarian conducts research in the computational areas of successive over-relaxation methods for large scale mathematical programs, normal solutions of mathematical programs, penalty functions in optimization, solution bounds for optimization

problems and condition numbers for linear and convex inequalities. Research emphasis is on the development of computationally oriented theory and effective algorithms for solution of large scale problems.

University of Wisconsin - Madison; Robert R. Meyer: *Parallel Computing for Large-Scale Optimization (Computer Research)*; (DCR-8503148); $45,943; 12 mos.

Most large-scale optimization problems exhibit special structures that allow the possibility of attack via algorithms that exhibit a high level of parallelism. The CRYSTAL multicomputer under development at the University of Wisconsin - Madison offers an opportunity to explore distributed approaches to optimization in that it provides a collection of loosely-coupled minicomputers of substantial computing power (specifically, VAX 11/750's) that can exchange data efficiently via messages. The emphasis of this research is on the development of optimization algorithms that take maximum advantage of the type of parallelism offered by loosely-coupled systems such as CRYSTAL. Approaches being investigated, including decomposition and branch-and-bound, are applicable to a wide variety of large-scale optimization problems and have the high granularity and relatively low data exchange that are optimal for distributed computing.

University of Wisconsin - Madison; Stephen Robinson; *Computational Approximation of Optimization Problems (Computer Research)*; (DCR-8502202); $43,986; 12 mos.

Professor Robinson is studying efficient methods for approximating optimization problems with applications to computational solution methods for large-scale problems encountered in resource allocation applications of both deterministic and stochastic types.

University of Wisconsin - Madison; Grace Wahba; *Multivariate and Multiresponse Estimation (Mathematical Sciences and Computer Research)*; (DMS-8404970 A01); $11,500; 12 mos; (Joint support with the Statistics and Probability Program - Total Grant $44,000).

This research is aimed at treating multivariate data where the problems are difficult, the need for good solutions is pressing, and the computational questions challenging.

Algorithms and Software

University of Arizona; Mike Szilagyi; *High Performance Electrostatic Lenses for Ion Beam Lithography*; (ECS-8317485 A01); $5,995; 12 mos; (Joint support with the Solid-State and Microstructures Engineering Program - Total Grant $59,947).

The use of charged particle beams in the semiconductor area is widespread and encompasses ion implantation techniques as well as surface studies such as SIMS. For future VLSI circuits, focussed ion beams may be used for direct writing of patterns and focussed ion droplets can be used for direct writing of metallic interconnects and mask repair. Present ion beam optics suffer deficiencies which appear as abberations limiting the focussed spot size and ion collection efficiency. This research addresses this problem by considering new ways to develop electrostatic lenses for application to ion beam lithography. Specifically the approach is computer based and utilizes approximation for the electrostatic potential distribution; the original electrode configuration can then be restructed from this. The approximation and calculation will be handled using the artificial intelligence language LISP in conjunction with an algebraic language MACSYMA. An expert system is being developed which will have the full capability of ion optical synthesis and electrode designs.

Stanford University; Walter Murray; *Computer Optimization of Complex Systems (Computer Research)*; (DCR-8413211); $205,000; 12 mos.

The objective of this research is to develop effective methods for solving large complex optimization problems. The areas of special concentration include nonlinear optimization algorithms, staircase (dynamic) linear programs, and linear and nonlinear complementarity problems. Work on sequential quadratic programming methods for nonlinearly constrained problems is proceeding in four main areas: construction of an augmented Lagrangian merit function for which global convergence can be proved, development of reliable methods for the case when second derivates are available, quasi-Newton updates of the projected Hessian of the Lagrangian function, and techniques for treating infeasible subproblems. For large-scale optimization problems, new techniques are being developed to compute and update a sparse factorization of the basis, subject to addition or deletion of rows or columns. Procedures are also being examined for dealing with partitioned systems in optimization algorithms, so that an initial factorization can be retained. Techniques based on dual nested decomposition of stochastic linear programs are being investigated, with special emphasis on making effective use of supercomputers and new computer architectures. For linear complementarity problems, a new multiparametric procedure is being studied in more detail, including investigation of its application to problems with degenerate matrices.

University of California - Los Angeles; N. Donald Ylvisaker; *Statistics and Game Theory (Mathematical Sciences and Computer Research)*; (DMS-8301587 A03); $20,400; 12 mos; (Joint support with the Statistics and Probability Program - Total Grant $86,950).

Research is being carried on in the general areas of statistics and game theory. Professor Jennrich is working on a number of problems related to his statistical computing interests. He is developing algorithms for both exploratory and confirmatory factor analysis and attack the problem of providing standard errors for maximum likelihood estimates in this context. Regression algorithms are being investigated as well. In the area of survival analysis he is investigating the behavior of tests for the comparison of survival curves under unequal censoring. Ylvisaker's research is concerned with model-robust design of experiments. Designs and estimates are being studied for the estimation of regression parameters in models which are infinite dimensional approximations to standard finite dimensional models. Corresponding work is underway for the robust modification of designs and estimates in combinatorial design settings. Professor Ferguson is working on stochastic games with information structure in which players have partial information about which game is being played. This area is a relatively unexplored one and he is considering questions of existence of value, form of optimal strategies, and the connections of this work with Bayesian statistics.

Yale University; John A. Hartigan; *Clustering Algorithms (Computer Research and Mathematical Science)*; (DCR-8401636 A01); $41,141; 12 mos; (Joint support with the Statistics & Probability Program - Total Grant $82,141).

The aim of the research is the development, evaluation, and distribution of clustering algorithms. New algorithms for block clustering and for reconstructing evolution from molecular data are being developed. Statistical evaluation takes data to be a random sample from a population, and considers how well various clustering techniques discover clusters in the population. The dip test of multimodality in one dimension is being studied further to establish its power against various alternatives. Extensions to many dimensions based on the minimum spanning tree are being explored. The asymptotic behavior of the minimum spanning tree (connected to single linkage clustering and percolation processes) is being studied. The block clustering algorithm BMDQ3M needs statistical tests to decide when the small blocks may be ignored.

Argonne National Laboratory; Jack Dongarra and Eric Grosse; *Mathematical Software Distribution via CSNET (Computer Research)*; (DCR -8419437); 36 mos; (Joint support with the Office of Advanced Scientific Computing - Total Grant $77,487).

Drs. Dongarra and Grosse are developing a system, using CSNET, to distribute mathematical software that is in the public domain. Such a system will provide for fast, easy-to-use, and inexpensive distribution of mathematical software.

University of Illinois-Urbana; Ahmed Sameh; *Parallel Numerical Algorithms*; (DCR-8117010 A03 & DCR-8509970); $35,110 + 10 hours of supercomputer time; 12 mos; (Joint support with the Office of Advanced Scientific Computing - Total Grant $70,110).

This research is developing numerical algorithms for solving important linear algebra problems on multiprocessors with two-level parallelism such as the CRAY-XMP48, GF10, and the Cedar. The research is considering the following:

1. Iterative schemes for solving positive definite linear systems, particularly those which arise from three-space dimensional parabolic problems;
2. Direct and iterative schemes for solving sparse linear least squares problems. These iterative schemes are also suitable for solving the linear systems which arise from certain non-self adjoint elliptic problems for which the available Chebyshev and C.G.-like schemes are not applicable; and
3. Solving the sparse symmetric and nonsymmetric generalized eigenvalue problem.

Massachusetts Institute of Technology; Virginia C. Klema; *Software Tools for Computation With IEEE Floating Point Arithmetic*; (DCR 8400246 A01); $102,611; 12 mos.

The recently approved IEEE binary standard for floating point arithmetic and its hardware implementation in microprocessors provide distinct enhancements for scientific computation. This research is directed toward the design and implementation of:

1. Software tools to access the floating point arithmetic exceptions; and
2. Basic linear algebra modules that optimize use of the floating point register stack.

These software tools are necessay to support the use of the binary standard for floating point arithmetic.

Massachusetts Institute of Technology; Roy E. Welsch; *Industry/University Cooperative Research Project: Nonlinear Exponential-Family and Bounded-Influence Regression (Computer Research)*; (DCR-8116778 A03); (Joint support with the Division of Industrial Science and Technology - Total Grant $37,947).

This is a collaborative research project between John Dennis of Rice University, David Gay of Bell Laboratories, and Roy Welsch of the Massachusetts Institute of Technology. The goal of the project is to exploit the

common structure of several nonlinear optimization problems.

Princeton University; John M. Mulvey; *Solving Generalized Networks With Convex Objectives (Computer Research)*; (DCR-8401098 A02); $7,891 + 10 hours of supercomputer time.

The two primary research objectives of this project are:

1. To integrate recent algorithmic advances in network optimization and nonlinear programming, and
2. To design efficient solution strategies for an important problem-class falling between these two areas.

The focus of attention is on nonlinear networks with generalized arcs (losses/gains) and non-separable convex objectives (NLGN). This model is capable of representing numerous significant real-world problems, including electrical power generation, stochastic networks, air-traffic routing, and financial cashflow.

This supplement provides support for expanding the research by providing real-world data from large scale applications and by motivating algorithms for solving ultra-large problems.

Carnegie Mellon University; Ravindran Kannan; *Computational Complexity of Numerical Algorithms (Computer Research)*; (DCR-8416190); $14,000; 20 mos; (Joint support with the Theoretical Computer Science Progam - Total Grant $67,430).

Improved algorithms for computing the Smith and Hermite normal forms of an integer matrix along with necessary empirical improvements are being studied. Theoretical improvements and implementations of algorithms for solving the same problems for matrices of polynomials are also being studied. The "alternation translation" technique is studied in order to obtain results about nondeterminism and determinism using the polynomial time hierarchy. Finally, several computational problems involving multivariate polynomials are being attacked.

University of Pittsburgh; Werner C. Rheinboldt; *Workshop on Future Directions in Computational Mathematics, Algorithms, and Scientific Software, Washington, D.C., November 20-21, 1984 (Mathematical Sciences and Computer Research)*; (DMS-8503483); $12,836; 12 mos; (Joint support with the Applied Mathematics Program - Total Grant $25,673).

A workshop to consider the current state of research and analyze future directions in computational mathematics, algorithms, and scientific software was held in the Fall of 1984. A panel of about twelve experts in the field assessed directions, established longer range objectives for the field, defined major subareas and identified typical research topics, analyzed possible forms of support needed and scrutinized current support efforts. The panel's findings are in the form of a report to the scientific community.

University of Texas - Austin; Frederick A. Matsen; *A Conference on Algorithms, Architectures, and the Future of Scientific Computation, Austin, Texas, March 18-20, 1985*; (DCR-8500515); (Joint support with the Office of Advanced Scientific Computation - Total Grant $9,000).

As new computers of wide range of capabilities are being considered for development as commercial products, there is an ever increasing need for developing strategies in the design of efficient algorithms and algorithm specific computers and in the promotion of large-scale computing. This conference considered such strategies.

University of Texas - Austin; David M. Young and David R. Kincaid; *Research And Development of Iterative Algorithms and Software for Elliptic Partial Differential Equations (Computer Research and Mathematical Science)*; (DCR-8214731 A02); $75,495; 12 mos; (Joint support with the Applied Mathematics Program -Total Grant $96,495).

This project involves research and development of iterative algorithms and software for solving large systems of linear algebraic equations with sparse matrices. Particular emphasis is placed on linear systems arising from the numerical solution of elliptic partial differential equations.

The work has two aspects, namely, development of research-oriented software and research on iterative algorithms. The first aspect of the work involves a substantial expansion of the capability of the ITPACK package to include more powerful iterative algorithms and to allow for handling more general classes of linear systems including systems with nonsymmetric matrices. It also involves continued participation in the ELLPACK project. The second aspect of the work is closely related to the first and involves research on iterative algorithms with emphasis on effective methods for solving nonsymmetric linear problems.

Symbolic and Algebraic Manipulation

Rensselaer Polytechnic Institute; Erich Kaltofen; *Complexity Studies in Computer Algebra (Computer Research)*; (DCR-8504391); $28,559; 12 mos.

Professor Kaltofen's research focuses on effective or efficient algorithms, both sequential and parallel, for solving problems in computational algebra. The problems to be addressed concern polynomial factorization, symbolic integration, number arithmetic, linear algebra, Galois theory, and cryptography.

Kent State University; Paul S. Wang; *Design and Implementation of Computer Algorithms for Symbolic Manipulation (Computer Research)*; (DCR-8504824); $34,000; 12 mos.

This research consists of the following four specific topics:

1. Extending current p-adic algorithms and programs to finite fields, rings and algebraic number fields.
2. Designing and implementing a much improved user interface for symbolic computation.
3. Developing techniques and applications of the symbolic derivation and generation of programs for numerical computation, and
4. Integrating graphics capabilities in symbolic computation systems.

This research should extend the frontiers of knowledge related to:

1. Efficient algorithms for exact symbolic mathematical computation involving p-adic algorithms;
2. The user interface for symbolic systems;
3. The use of symbolic computation to prepare programs for numerical calculations; and
4. The application of interactive graphics in symbolic computation.

Southern Methodist University; David Y. Yun; *Algorithms and Systems For Symbolic Algebraic Computation (Computer Research)*; (DCR-8314600 A03 & A04); $96,543; 12 mos; (Joint support with the Industry/University Cooperative Research Projects Program - Total Grant $193,043).

This project involves the study of algorithms for symbolic and algebraic computation. The effort embraces pure mathematical and algorithmic research as well as implementation and testing of algorithms with the goal of achieving efficiency and practicality. The research is being conducted by B. David Saunders and S. K. Abdali in the Department of Mathematical Sciences at Rensselaer Polytechnic Institute, Richard D. Jenks and Barry M. Trager in the Mathematical Sciences Department of IBM Research, Yorktown Heights, New York, and David Y. Yun of the Department of Computer Science and Engineering at Southern Methodist University.

The objectives of the research program are to discover, analyze, implement, and test algorithms related to the automatic solution of equations in closed form.

The research program consists of four parts:

1. Investigation of algorithms for solving Diophantine equations over unique factorization domains,
2. The development of a subsystem in computer algebra for linear algebraic computations over non-numerical fields,
3. Investigation of efficient algorithms for solving systems of polynomial equations, and
4. Study of Liouvillian theories and development of algorighms for solving linear ordinary differential equations.

University of Washington; Martin Tompa; *The Combinatorial Structure of Computations and Symbolic Manipulation (Computer Research)*; (DCR-8301212 A02); $15,392; 12 mos; (Joint support with the Theoretical Computer Science Program - Total Grant $30,785).

The success of theoretical computer science is due largely to the process of abstracting clean, tractable problems from the morass of realistic detail. Yet there is much left to be desired in the understanding of the structure of computations. For instance, it is not known how to exploit highly parallel systems. Neither is it known how to exploit randomness (in the sense of probabilistic, or Monte Carlo, computation) to derive efficient algorithms. Even in the setting of ordinary sequential computation, where a reasonably good repertoire of algorithms exists, there are few techniques to show that those algorithms are optimal.

Part of the problem is that not enough of the detail is abstracted away. General computational models like Turing machines or random access machines have sufficient detail to defy analysis.

In this project, computations are abstracted further in order to lay bare the underlying combinatorial structures. With all inessential details stripped away, it is often simpler to discover new efficient algorithms, and demonstrate their optimality (at least within this structured framework) than it is with the details included. The areas being studied include synchronous parallel complexity, sequential time and space complexity, and probabilistic complexity.

University of Wisconsin - Madison; George E. Collins; *Algebraic Algorithms and Systems (Computer Research)*; (DCR-8408107 A01); $73,648; 12 mos.

This project involves the study of five related research areas in symbolic and algebraic manipulation. These topics are:

1. Polynomial factorization,
2. Cylindrical algebraic decomposition,
3. Polynomial residue class rings,
4. Linear diophantine equations, and
5. Development of the SAC-2 symbolic manipulation system.

Computer Systems Design

VLSI Design Methodology

University of Southern California: Melvin Breuer and Alice Parker; *The Adam VLSI Design System*; (DMC-8310774 A01); $35,000; 12 mos; (Joint support with the Computer Engineering Program - Total Grant $140,000).

This project addresses that research essential to build an Advanced Design Automation System (ADAM). ADAM is an experimental system capable of being used for the design, and implementation in silicon, of complex digital computing systems.

Three major research issues are being explored, namely:

1. Overall system architecture of ADAM,
2. The design of a data base system to support ADAM, and
3. Problem specific applications such as synthesis, verification, layout and testing.

This research covers the first two major issues and a single application domain; an expert system for custom layout (CLOUT). CLOUT is being developed as part of the overall design system because physical design (including layout) proceeds in parallel with logical design. Thus to evaluate CLOUT, one needs a total design environment with feedback on system architecture, choice of layout design styles and tradeoffs in performance.

There are three novel aspects to the work:

1. The emphasis on embedding artificial intelligence techniques into a general computer-aided design automation system architecture, not just for a particular application,
2. The design data being structured in a new way which separates target system behavior, structure, control and timing, and physical information, and
3. Much of the design data and many of the steps in the design process being characterized mathematically which provides a continuum between design and analysis.

Washington University: Mark A. Franklin and Donald F. Wann; *Logic Simulation Algorithm and Machine Design (Computer Research)*; (DCR-8417709); $64,974; 12 mos; (Joint support with the Computer Engineering Program - Total Grant $129,948).

Increased circuit densities resulting in high total circuit element counts have made logic simulation a critical problem in the design and testing of VLSI-based chips. Simulation times for large circuits are now often measured in days and even months. Although several specialized logic simulators have been constructed, there is a paucity of analysis of the complete simulation task in which the underlying fundamental issues are explored and then exploited in the design of hardware-based logic simulators.

Cornell University: Kevin Karplus; *Clocking Disciplines and Mutual Exclusion Constraints Formal Models and Verification Tools (Computer Research)*; (DCR-8503262); $104,901; 24 mos.

This project focuses on the study of the theoretical underpinnings of clocking disciplines and mutual exclusion constraint rules, and the incorporation of the results of these studies into CAD tools. Clocking disciplines are an approach for designing synchronous systems that have no race conditions. Mutual exclusion constraints are boolean equations that must always be satisfied by the signals in a circuit for the circuit to be adequately represented by simple switch models.

Carnegie-Mellon University: Stephen W. Director, Daniel P. Siewiorek, and Donald Thomas; *Multilevel Computer-Aided Design of VLSI Digital Systems*; (DMC-8207709 A02); $101,433; 12 mos; (Joint support with the Electrical and Optical Communications Program - Total Grant $202,867).

A hierarchial approach to the design of digital systems is being taken and begins with a behavioral specification of the system. This specification provides a model which accurately characterizes the input-output behavior of the system without reflecting any internal structure. The hierarchial design is supported by a comprehensive set of computer aided design (CAD) methods which takes into account the important characteristics of integrated circuit technology.

This research continues previous work and addresses the following areas:

1. Synthesis of digital VLSI sytems,
2. Multilevel representations for digital VLSI systems,
3. Multilevel optimization and digital system design space exploration,
4. Man-machine environment for CAD of digital VLSI systems, and
5. Special hardware for VLSI CAD.

Defense Advanced Research Projects Agency; Robert S. Cooper; *NSF/DARPA Agreement for Use of DARPA VLSI Implementation Service*; (DMC-8543305); $100,000; 12 mos; (Joint support with the Computer Engineering Program - Total Grant $200,000).

The development of technology making feasible Very Large Scale Integrated (VLSI) chips opens the way for universities to play a major innovative role in the design, construction, and application of novel special purpose hardware and new computer architectures. Considerable research is already in progress in the university community on software and hardware design tools needed to manage the complexity which arises in designing new VLSI computing structures containing tens of thousands of transistors on a silicon chip. Novel designs are being proposed and studied. There is a great need to implement these designs in silicon chips and test their performance.

As part of DARPA's research in the VLSI area they have developed a capability called "MOSIS" for fast turnaround implementation of VLSI chips. Designs submitted by DARPA researchers in digital form over the ARPANET or Telenet using standard design rules and a standard artwork format, known as CIF, are fabricated and returned to designers in about 4-6 weeks.

This represents the fourth and final year funding of the four year joint experiment being undertaken by the National Science Foundation and the Defense Advanced Research Projects Agency (DARPA) that allows qualifying universities to use the DARPA fast turnaround VLSI implementation facility for university-based research and educational programs requiring the fabrication of digital designs as integrated circuit chips.

Access to this VLSI service by universities is on a controlled basis with usage limited to those institutions that meet requirements jointly established by NSF and DARPA. The integrated circuit chips are fabricated, packaged, and sent to the university groups submitting the designs, as soon as possible.

Computer System Architecture

Arizona State University; Antony Faustini; *An Eduction Engine for Lucid - A Supercomputer Based Upon Demand-Driven and Data-Driven Evaluation (Computer Research)*; (DCR-8415618); $56,774; 12 mos.

A new approach to computer architecture, based on the use of demand-driven asynchronous computation, the use of tagged data to obviate the need for 'state' information, and the use of an existing high level language (Lucid) as the machine language is being explored. The approach is to design an experimental eduction engine, i.e., a highly parallel machine that uses tagged data with the demand-driven approach to asynchronous distributed computation. Analysis of the behavior of the machine and the behavior of Lucid programs is aimed at realizing a truly scalable machine which is not restricted to special applications or to special types of algorithms. The effects of different design decisions are being explored via extensive simulations. These simulations are also being used to confirm the expected machine performance.

This is a collaborative project with Edward Ashcroft at SRI International.

SRI International; Edward Ashcroft; *An Eduction Engine for Lucid - A Supercomputer Based upon Demand-driven and Data-driven Evaluation (Computer Research)*; (DCR-8413727); $73,490; 12 mos.

A new approach to computer architecture, based on the use of demand-driven asynchronous computation, the use of tagged data to obviate the need for 'state' information, and the use of an existing high level language (Lucid) as the machine language is being explored. The approach is to design an experimental eduction engine, i.e., a highly parallel machine that uses tagged data with the demand-driven approach to asynchronous distributed computation. Analysis of the behavior of the machine and the behavior of Lucid programs is aimed at realizing a truly scalable machine which is not restricted to special applications or to special types of algorithms. The effects of different design decisions are being explored via extensive simulations. These simulations are also being used to confirm the expected machine performance.

This is a collaborative project with Antony Faustini at Arizona State University.

University of California - Berkeley; Alvin M. Despain and Vason P. Srini; *Dynamically Reconfigurable Operating Systems Research (Computer Research)*; (DCR-8508344); $40,140; 15 mos.

This project focuses on the development of dynamically reconfigurable operating systems (DROS) for distributed systems using a dataflow methodology. DROS refers to changing the strategies and the policies (semantic change) used by the modules of a distributed operating system in a dynamic manner. Dynamic semantic changes to modules of an operating system (OS) are important in distributed systems. This allows the OS to use different algorithms for task assignment,

memory management, and interprocess communication as the system resources and workload change.

One goal of the research is the development of a methodology for the design of a DROS using dataflow graphs. The methodology supports modular design, stepwise refinement, interprocess communication using self identifying messages (datagrams), and reconfiguration. Another goal is the development of reconfiguration strategies based on global information, local information, and a combination of the two. Techniques for adapting the reconfiguration strategies to satisfy data dependency and realtime needs are also being developed.

University of Illinois - Urbana; David J. Kuck and Duncan H. Lawrie; *Cedar Systems Prototype*; (DCR-8410110); $266,522; 12 mos.

The primary goal of the Cedar project is to demonstrate that supercomputers of the future can exhibit general-purpose behavior and be easy to use. A major prototype is being constructed to demonstrate the significance of the Cedar architecture and to further explore the underlying concepts. The following problem areas are being addressed:

1. Communications in a large-scale multiprocessor,
2. Global control of a large-scale multiprocessor,
3. The synchronization of processes, processors, tasks, and computations,
4. Compilation for a vector multiprocessor,
5. Compile-time memory hierarchy management,
6. Multiprocessor operating systems, and
7. Algorithms for systems with concurrency.

Louisiana State University - Baton Rouge; Walter G. Rudd; *The Factoring Machine (Computer Research)*; (DCR-8506191); $40,261; 12 mos.

General purpose computers are not well suited for the long-wordlength, highly-parallel, fast integer arithmetic that factoring, primality testing and other similar problems in experimental number theory require. This project focuses on the design and construction of the Factoring Machine, its testing using existing algorithms for factoring large numbers, the development of new algorithms that take advantage of the machine's architectural features, and the investigation of the problem of software development for machines with this structure.

University of Michigan - Ann Arbor; Arthur W. Burks; *Languages and Architectures for Parallel Computing with Classifier Systems (Computer Research)*; (DCR-8305830 A02); $105,042; 12 mos; (Joint support with the Intelligent Systems Program - Total Grant $210,084).

A new class of rule-based computing systems, classifier systems, has been developed and tested. While classifier systems are similar to production systems, they are different in basic ways which facilitate the incorporation of learning algorithms in them, and which make it possible for them to be executed in a highly parallel fashion without the usual problems of concurrency and interlock. Classifier systems with learning algorithms have been tested successfully on two model control systems, a game learning program, and a hardware robot project. Progress has been made on the theoretical understanding of how such classifier systems work.

This research focuses on two interrelated areas:

1. The powers and limits of classifier systems are being studied by further computer simulations and theoretical analyses; and
2. Novel parallel architectures for the rapid execution of classifier systems are being simulated.

On the theoretical side, these studies contribute to a deeper understanding of concurrency and distributed control. On the practical side they lead ultimately to powerful parallel computers which are non-von Neumann in both programming and architecture.

Cornell University; Jon Solworth; *Negotiators: A Novel Architecture for Parallelism (Computer Research)*; (DCR -8503610); $106,822; 24 mos.

A novel architectural structure called Negotiators which replaces conventional memory banks with "smart memory" and which supports classes, monitors, and very high level languages is being explored. A simulator is being built and a compiler modified for the purpose of evaluating Negotiators. Benchmark programs are being obtained and written to measure performance improvements. These benchmarks provide the basis on which to parameterize the design. Compiler suupport for such structures is also being explored.

New York University; Jacob T. Schwartz and Allan J. Gottlieb; *VLSI Architecture for Ultraparallel Computers (Computer Research)*; (DCR-8413359); $260,000; 12 mos.

The primary characteristic that distinguishes the NYU Ultracomputer from other shared memory MIMD parallel computers is its enhanced Omega-network, capable of combining simultaneous references to the same memory location. This project focuses on two approaches to the design of this combining network. The first involves a crossbar arrangement of systolic queues. To date, a VLSI version of a systolic queue has been designed, constructed, and tested. A 2x2 (buffered) crossbar has been constructed and tested from these systolic queues. An 8x8 (bidirectional) Omega-network is being constructed and incorporated in the currently operational Ultracomputer prototype in order to further explore the effectiveness of this type of interconnection network. The second form of interconnection network being explored consists of a

collection of INMOS Transputers, each containing a simulator of the combining switch. Use of this interconnection approach allows the investigators to explore alternative interconnection strategies by simply modifying the combining strategies in the simulator. The current NYU Ultracomputer design appears attractive for configurations ranging from hundreds to tens of thousands of processors. This question of scaleability is the focus of the third major element of this project.

Southern Methodist University; David W. Matula; *Foundations of Finite Precision Rational Arithmetic (Computer Research)*, (DCR-8315289 A01); $40,167; 12 mos.

VLSI capabilities are fostering an era wherein complex special purpose computer hardware for specific broad application areas is realizable at acceptable cost. Previous research by Dr. Matula has described special number systems and proposed a new computer arithmetic unit sufficiently competitive in convenience of arithmetic logic design and timing efficiencies to be a serious candidate for a special purpose arithmetic architecture. This research project is further refining the design possibilities for finite precision rational arithmetic by:

1. Investigating the utilization of binary continued fraction representation for values in storage and as input/output of a rational arithmetic unit,
2. Incorporating on-line techniques in chain rational computation and other areas for increased parallelism in a rational arithmetic unit,
3. Utilizing extended range binary floating-slash representation and/or binary continued fraction representation to prescribe a hierarchical precision rational computation system, and
4. Investigating accumulated roundoff error growth for computations hosted in alternative finite precision rational number systems.

University of Utah; Lee Hollaar; *Continued Analysis and Development of Specialized Backend Processors for Very Large Text Databases (Computer Research)*; (DCR -8503663); $156,864; 24 mos.

The Utah Text Search Engine is a special-purpose backend processor capable of scanning serial data from a mass storage device for occurrences of complex patterns. It is based on a new form of finite state automaton, the partitioned FSA, which is well suited for VLSI implementation. Custom integrated circuits have been developed and tested, and an initial prototype configuration has been used to replace the software-based search module in the prototype retrieval system runnning on the University of Utah Computer Science Department Apollo workstation network. A new prototype is currently under development.

While the prototypes have demonstrated the feasibility of the PFSA-based search technique, a detailed evaluation of its performance in a large system environment has not been made. The effects on its overall performance for different system configurations and sizes, as well as different index strategies, are being simulated and analyzed. State table characteristics for different query constructs, such as don't cares, are being examined. Finally, extensions to the basic PFSA technique to handle non-text databases are being explored.

Virginia Polytechnic Institute & State University; James Armstrong; *Chip Level Modeling of Digital Systems (Computer Research)*; (DCR-8420687); $113,046; 24 mos.

The advent of Very Large Scale Integration (VLSI) technology has rendered the gate level model impractical for many modeling activities critical to the design automation process. As an alternative, an approach to the modeling of VLSI devices at the chip level is proposed which will allow the chip to be the basic logic primitive in systems built from a large number of VLSI devices. Work on this concept was begun under previous research grants and has resulted in a fully implemented chip level modeling and simulation system. This effort is concentrated on developing an effective system for modeling faults at the chip level. The goal is the development of an automatic system which could inject chip level faults and determine tests for these faults. Another effort involves improving the efficiency of chip level simulation through the development of a parallel computer architecture for this application.

University of Washington; Jean-Loup Baer; *Cache Coherence in Multiprocessor Systems (Computer Research)*; (DCR-8503250); $128,709; 24 mos.

The management of the memory hierarchy in tightly-coupled multiprocessor systems with particular emphasis on the cache coherence problem is being explored. The main areas of investigation are:

1. The implementation and evaluation of protocols and an associated memory controller for a cache coherence solution that has been proposed recently;
2. A comparative evaluation of various cache coherence solutions for shared bus systems;
3. Cache coherence solutions for large systems with a particular emphasis on systems with cache hierarchies;
4. Trade-offs between conventionally addressed systems and those where caches are addressed via virtual addresses; and
5. Application of similar mechanisms for data coherence in non-transactional applications in networks.

University of Wisconsin - Madison; Leonard Uhr; *Algorithm-Structured Pyramid and Network Architectures (Computer Research)*; (DCR-8302397 A02); $73,185; 12 mos.

This research continues an on-going investigation of multi-computer networks, in an attempt to find a) good algorithm-structured architectures, b) good general network architectures, and c) good criteria, both structural and formal, for evaluating specialized and general multi-computer architectures.

Several key issues are being examined related to future VLSI implementation of extremely large multi-computers. These include:

1. How are basic resources best distributed among processors, memories, reconfiguring switches and controllers?
2. How are micro-modular arrays and pyramids best designed?
3. How can fault-tolerant reconfiguring capabilities that take advantage of the micro-modular array pyramid structure best be added?

Fault Tolerance and Reliability

Stanford University; Edward J. McCluskey; *Reliable Computers (Computer Research)*; (DCR-8200129 A03); $109,627; 12 mos.

This research is aimed at developing systematic design techniques for high-availability computer design. Specific attention is being given to unification of the treatment of:

1. Production testing of integrated circuits and inclusion of appropriate circuitry to guarantee that this is economical (Design for Testability),
2. Circuitry for concurrent checking of computer operation, and
3. System structure to facilitate restructuring to eliminate failed subsystems from interfering with correct operation.

The novelty of this study stems from the fact that all the aspects of high-availability computer design will be treated in a uniform way. Past studies of this topic have tended to emphasize one aspect of the problem to the exclusion of the others. Modern VLSI techniques make it possible to combine the various techniques into one unified design technique.

University of California - Los Angeles; David A. Rennels and Algirdas Avizienis; *Fault-Tolerant Local Networks Using VLSI Building Blocks (Computer Research)*; (DCR-8307026 A01); $114,216; 12 mos.

This research is aimed at developing and evaluating architectural concepts for fault-tolerant local networks using VLSI building block circuits to provide hardware fault-tolerance and a software environment characterized by a high level of security, ease of programming, and run-time checking for a wide range of software errors.

The major areas of study are:

1. An investigation of architectures for large shared memory and computational resources which are self-managing and fault tolerant,
2. An investigation of the applicability of specialized hardware support in intercommunications and synchronization, and
3. A study of smart memory interfaces which automatically map the local address spaces of distributed computers to the virtual address space provided by such a network.

University of Michigan - Ann Arbor; John P. Hayes; *Hierarchical Testing of Digital Systems (Computer Research)*; (DCR-8419989); $139,953; 24 mos.

Efficient test generation algorithms and self-testing design methods for complex digital circuits are being explored. A hierarchical approach spanning the switch, gate, and register complexity levels is employed. The CSA (connector-switch attenuator) circuit model and a very general class of logical fault models is being used, along with a new technique for bidirectional signal evaluation and error propagation. Trade-offs between fault coverage, modeling level, and computational complexity are being studied. A novel self-testing scheme that has been successfully applied to the design of a dynamic RAM chip is being extended to more general circuits. It introduces self-testing by partitioning a circuit in hierarchical fashion into (nearly) identical subcircuits that can be tested in parallel. The resulting responses are verified by on-chip comparators. Also being investigated is the problem of obtaining quantitative measures of the self-testability of a design.

Computer System Performance Measurement and Evaluation

University of Massachusetts - Amherst: Victor R. Lesser: *Coordination in Cooperative Distributed Problem Solving Systems (Computer Research)*; (DCR-8300239 A02); $42,348; 12 mos; (Joint support with the Intelligent Systems Program and the Special Projects Program - Total Grant $128,696).

A key problem in cooperative distributed problem solving is developing network coordination policies that provide sufficient global coherence for effective cooperation. This problem is difficult because limited internode communication precludes the use of a global "controller" node. Instead, each node must be able to direct its own activities in concert with other nodes, based on incomplete, inaccurate and inconsistent information. This requires a node to make sophisticated local decisions that balance its own perceptions of appropriate problem solving activity with activities deemed important by other nodes. Current research at the University of Massachusetts has included the development of a node architecture capable of such sophisticated local decision making. This architecture has been implemented as part of the distributed vehicle monitoring testbed: a flexible and fully instrumented research tool for the empirical evaluation of distributed network designs and coordination policies. This research builds on this node architecture and testbed to explore (through actual implementation and empirical studies) a variety of approaches to network coordination that include: organizational self-design, distributed load-balancing, negotiation among nodes, planning of internode communication strategies, and knowledge-based fault-tolerance.

North Carolina State University: Harry Perros: *Networks of Queues with Blocking (Computer Research)*; (DCR-8502540); $95,000; 24 mos.

Queueing networks with blocking frequently arise in performance studies of computer systems, distributed systems and telecommunication systems. Despite their potential applicability their use is not widespread. This is mainly due to the absence of comprehensive tools for analyzing such queueing systems.

This research is focused on the following three tasks:

1. Devise exact and approximation procedures for analyzing arbitrary configurations of queueing networks with blocking,
2. Incorporate these procedures into an existing algorithm for analyzing general queueing networks, and
3. Carry out a comprehensive survey of the literature on queueing networks with blocking.

North Carolina State University - Raleigh: William Stewart; *Approximate Solutions of Queueing Models — An Error Analysis (Computer Research)*; (DCR-8503067); $15,000; 12 mos; (Joint support with the Software Engineering Program - Total Grant $45,881).

This research investigates the error behavior of iterative approximation methods. Phase I characterizes iterative methods as systems of non-linear equations and also categorizes the problems to which the iterative methods have been applied to develop a comprehensive testbed of models. Phase II is a systematic study of the errors generated by iterative methods. The study involves implementing the different methods and generating results for the models in the testbed over a range of parameter values. Mathematical expressions for the errors are being developed.

New York University: Uzi Vishkin: *Synchronous Parallel Computation-Efficiency of Resources (Computer Research)*; (DCR-8318874 A01); $22,700; 12 mos.

Future ultra-large computers will be built with a design space (i.e., the model for which programs are written) which is very permissive. A certain permissive model of parallel computation is studied. The implementation space must meet the limitations of available technologies. The interrelation between the design space and several suggestions for both the implementation space and simulation of the design space into the implementation space is being analyzed. Algorithms in the design space concentrating on the complexity of the time, number of processors, and communication facility are being studied.

Computer Graphics

Princeton University: David P. Dobkin: *Theoretical and Pragmatic Concerns in Computer Graphics (Computer Research)*; (DCR-8303926 A02); $101,162; 12 mos.

This research is exploring the implementation of graphics primitives. The study is at the theoretical level but also involves hardware and software implementa-

tions. Three problems are being explored in detail. These represent different parts of the graphics process - line drawing, hidden surface elimination and number representations for accurate computation.

Cornell University; Donald P. Greenberg; *Interactive Computer Graphics Input and Display Techniques (Computer Research)*; (DCR-8203979 A03); $252,832; 12 mos.

This research project is concerned with the development of new and improved input and display methods for interactive computer graphics. Two general areas are addressed, the graphical creation of geometric object descriptions, and algorithmic processes for realistic and dynamic image synthesis. Specific research topics include color science and equal perception spaces, a polygon hidden surface algorithm, a light reflection model, and investigations into new uses of coherence for high quality image display. Results are potentially useful to a full range of computer-aided-design applications.

Ohio State University; Charles A. Csuri and Edwin Tripp; *Motion Control, Natural Phenomena, and Parallel Architectures for Display Algorithms (Computer Research)*; (DCR-8304185 A02); $210,000; 12 mos.

Three areas of research are being investigated. Primary emphasis is on the areas of motion control and parallel architectures for display architectures. Secondary emphasis is on the simulation of natural phenomena. There is no simple or short term solution to any of these tasks. The control of articulated motion presents problems of algorithmic control, knowledge-based animation, movement knowledge base, environmental knowledge base, special problems in kinematics, the application of special purpose motion processors and the generation of complex surfaces for moving figures. Parallel architectures for display algorithms, for the most part, have been described at a theoretical level. There are few case studies to validate various theories. An experimental parallel microprocessor architecture is available on which initial studies on the partitioning of display algorithms are being explored. Simulation of natural phenomena such as clouds, fire, smoke, ice and even trees have not yet been displayed realistically in computer graphics. The problem is mainly concerned with data complexity and computational requirements.

University of Utah; Robert E. Barnhill; *Surfaces in Computer Aided Geometric Design (Computer Research)*; (DCR-8502858); $132,431; 12 mos; (Joint support with the Software Engineering Program - Total Grant $177,431).

This research deals with the creation of new surface forms and the determination of geometric properties of surfaces in an interactive computer graphics environment. The investigations of new surface forms include the topics of Bernstein-Bezier triangular interpolants, surfaces defined on surfaces, and sweep surfaces. The geometric properties considered are estimation of derivative data, visual (or geometric) continuity and reflection lines of surfaces. In creating surfaces, the geometric data given are "shape" information given either as "hard" data such as positions or as "soft" data such as a designer's conception of a pleasing surface. In typical problems one encounters a mixture of both types of data. The surfaces to be created must therefore satisfy a variety of criteria such as sufficient smoothness, reasonable shape, and sufficient accuracy or fidelity. These requirements cannot all be met by the standard tensor product surfaces often used in practice. This lack is due to the rigid rectilinear structure required by tensor product schemes. Consequently, the research takes the more ambitious approach of developing a robust class of surfaces that can also utilize arbitrarily located information. The research involves the mathematics of surfaces, algorithms for the rendering of surfaces, the computer graphics implementation of these algorithms, and the recording of the visual results by means of some photographic medium. This research area is widely applicable in "CAD/CAM" (Computer Aided Design/Computer Aided Manufacturing), an engineering discipline dedicated to the automation of design and manufacturing.

University of Washington; Kenneth R. Sloan, Jr.; *Reconstruction of the Surface of Three Dimensional Objects from a Collection of Planar Cross-section Contours (Computer Research)*; (DCR-8505713); $111,639; 24 mos.

This research project focuses on the reconstruction of three-dimensional surfaces from a collection of planar cross-sectional contours. Current solutions to this problem have raised interesting theoretical questions about search techniques and the exploitation of domain-specific aspects of such search problems. Existing implementations of these solutions fall into two broad classes: Those which use a "greedy" search to find an approximate solution, and those which use an exhaustive search to find an exact solution.

This research involves the building of a testbed for the investigation of the following questions:

1. How can exact solutions be found more rapidly?
2. How much better are the exact solutions than the approximate solutions?
3. Can the cases in which approximate solutions are sufficient be characterized?

Research also focuses on the evaluation of existing methods and on the analysis of these methods with a view towards discovering new techniques.

BLANK PAGE

Intelligent Systems

Computer Vision and Image Processing

Stanford University: Paul Switzer: *Statistical Theory and Methods for Processing Spatial Imagery (Mathematical Sciences and Computer Research)*; (DMS-8411300 A01); $14,000; 12 mos; (Joint support with the Statistics and Probability Program - Total Grant $42,500).

The research aims to develop models and algorithms to treat very complex and poorly understood issues connected with reconstructing images from noisy data.

University of California - Berkeley: Alberto Grunbaum: *Reconstruction with Limited and Noisy Data (Mathematical Sciences and Computer Research)*; (DMS-8403232 A01); $8,000, 12 mos; (Joint support with the Applied Mathematics Program - Total Grant $64,054).

The central issue in the data inversion problems being studied is to determine specific properties of an object from measurements of the way it transmits or reflects acoustic signals, thermal waves, X-rays or other types of incident waves. When the measurements are completely accurate and sufficiently abundant, the desired reconstruction can often be accomplished by elementary means. But in practice the data are often incomplete and corrupted by noise and reconstruction techniques which work for accurate and complete data may fail completely. This research is concerned with the reconstruction problem in such an environment. In particular, the problems of how best to adapt or modify algorithms designed to handle complete and clear data to the case of incomplete and noisy measurements are being investigated.

University of Miami: Tzay Y. Young: *Three-Dimensional Motion Analysis Using Shape Change Information (Computer Research)*; (DCR-8509737); $38,997; 12 mos.

This research is investigating methods for estimating 3D motion parameters by analyzing 2D shape changes of corresponding regions in consecutive images. Shape change analysis utilizes effectively regional motion information, and hence the estimation of motion parameters could be more accurate than the existing point correspondence approaches. Accurate identification of a number of corresponding points is difficult in most real images.

A general model of shape changes is being explored using Taylor series expansion of a 2D velocity field. Linear and nonlinear shape change parameters are defined, and parameter estimation algorithms are being developed. Quantitative relationships between these parameters and 3D motion are being established so that motion parameters can be computed. A model-based hypothesize-and-verify method that takes into consideration the nonlinear effect of perspective projection is being investigated. Software programs are being developed, tested and implemented. Experiments are being carried out using real images, and the sensitivity and accuracy of the proposed methods are being evaluated.

University of Illinois - Urbana: Thomas S. Huang: *Acquisition, Representation, and Manipulation of Time-Varying Spatial Information (Computer Research)*; (DCR-8415325); $65,797; 12 mos.

The main motivation is the use of visual feedback in robotics. In the case of automated manufacturing, a robot is required to inspect parts for identification and diagnosis and to navigate for arranging objects and moving around on factory floors. All these tasks require perception and manipulation in the 3-D environment. This research is concerned with the analysis and modeling of time-varying spatial information focussing on three interrelated areas:

1. Matching and registration techniques,
2. Volume representation of time-varying spatial information, and
3. Range and motion determination by temporal and spatial stereo methods.

Purdue University: Avinash Kak: *Third Workshop on Computer Vision: Representation and Control*, October 13-16, 1985, Shanty Creek, Michigan; (DMC-8511216); $9,000; 12 mos; (Joint support with the Automation and Systems Integration Program - Total Grant $26,882).

This grant supports the third in a series of workshops on representation and control in computer vision. A set of position papers is being produced which outline the state of the art in computer vision and delineate reserach aspects that are properly the domain of industrial research laboratories. The focus is on fundamental research issues in computer vision with relevance to robotics and automation.

University of Maryland - College Park: Azriel Rosenfeld: *Multiresolution Image Analysis (Computer Research)*; (DCR-8218408 A03); $184,845; 12 mos.

During the past few years there has been increasing interest in the use of multiresolution ("pyramid") image representations for a variety of image processing and analysis tasks. This approach has two principal advantages. First, it often leads to reduced computational costs through the hierarchical structuring of computations (e.g., "divide-and-conquer"). Second, it has suggested some interesting new types of algorithms, involving cooperation among pixels or features at different resolutions, which may have advantages over conventional approaches, in particular by providing a smoother transition between pixel-level and region-level image representations. This research on multiresolution image analysis techniques is concerned both with the development of new methods and with their incorporation into an integrated image analysis system which allows pixel-level and region-level knowledge to interact through the medium of multiresolution representations.

This grant supports the summer research of Dr. Robert Melter, a small college faculty member at Southampton College, Long Island University.

University of Maryland - College Park, Hanan Samet: *Hierarchical Data Structures For Image Processing (Computer Research and Information Science and Technology)*; (DCR-8302118 A01); $31,658; 12 mos.; (Joint support with the Division of Information Science and Technology - Total Grant $61,658).

Hierarchical data structures such as quadtrees and octrees are becoming increasingly important representations in the domains of computer graphics, image processing, pattern recognition, computational geometry, and geographic information systems. The following areas are being investigated:

1. Development of algorithms for multi-dimensional data;
2. Incorporation of the time dimension into hierarchical representations,
3. Testing image approximation methods that are based on quadtrees;
4. Alternative recursive decomposition techniques and their applicability in a multiprocessor environment; and
5. Hierarchical representations of curvilinear data.

Massachusetts Institute of Technology: Whitman Richards: *Natural Computation: A Computational Approach to Visual Information Processing (Information Science and Computer Research)*; (IST-8312240 A01); $60,305; 12 mos. (Joint support with the Division of Information Science and Technology - Total Grant $180,916).

Natural computation is the study of how man can derive reliable representations of useful aspects of his environment. These representations may be visual, auditory, or consist of programs for touch sensing and motor control. At present, the data under study are the vast arrays of image intensities cast upon the eye or camera. Given these data, the task is to derive rather simple but reliable descriptions of objects and events in a three-dimensional world. The problems involve choice of representation, constraint analysis and explicit implementation (algorithm) to show that reliable and efficient descriptions can be computed.

The research may be divided roughly in four general problem areas:

1. 3D surfaces reconstruction,
2. Property-based representations,
3. Integrating depth maps obtained from stereo, structure-from-motion, and shading; and
4. Computing spatial relations between "parts" and their combination into "object" representations.

University of Massachusetts - Amherst: Edward M. Riseman and Arthur S. Gaylord: *A Group Research Facility for Artificial Intelligence, Distributed Computing, and Software Systems (Computer Research)*; (DCR-8318776 A02); $25,000; 12 mos; (Joint support with the Special Projects, and the Software Engineering Program - Total Grant $120,000).

Experimental computer research is focused in six areas: computer vision, robotics, distributed computing, software development, intelligent interfaces, and natural language processing. While most of the actual research activities are supported under other grants and contracts, including some from NSF, this grant provides partial support for the experimental research facility used by all the researchers, supplementing institutional, departmental, and other funds. NSF support is used for administration and operation of the COINS research facility, with the funding going towards technical support staff and equipment maintenance.

Wayne State University: Ishwar K. Sethi: *Trajectory Based Framework for Dynamic Scene Analysis (Computer Research)*; (DCR-8500717); $40,677; 12 mos.

This research is concerned with two main aspects of dynamic scene analysis: to develop techniques for feature point correspondence and three-dimensional structure perception by integrating the information from a long sequence of image frames. The key ideas behind the correspondence algorithms are the notions of path coherence and the smoothness of motion of physical objects. Following these two notions the goal is to develop techniques for corrrespondence which ensure the overall minimum deviation of the feature points in space-time. Methods leading to the closed form solution of the three-dimensional structure of the

moving objects are being developed based on the processing of feature point trajectories. Some of the related issues like the detection of moving feature points, and the problem of curve fitting with respect to highly eccentric ellipses are also being studied.

University of Minnesota: William B. Thompson: *Determining Spatial Organization from Visual Motion (Computer Research)*; (DCR-8500899); $39,697; 12 mos.

Visual motion is a powerful cue for specifying information about the spatial organization of a scene. Significant progress has been made in estimating optical flow and in deriving motion-based methods for determining properties of surfaces such as orientation. This research deals with techniques for finding spatial relationships between different surfaces. Improved methods for finding surface boundaries based on optical flow are being explored. The aim is to discover the types of information available at dynamic occlusion boundaries and to create effective algorithms for extracting that information. This work goes well beyond the simple detection of boundaries, as the pattern of visual motion at a boundary says much about the structure of the surfaces making up the boundary. Methods are being developed for determining the occluding and occluded surfaces generating a boundary. In turn, this leads to techniques for determining relative depth at boundaries. The ultimate goal is to develop algorithms for deriving an extensive set of constraints on the spatial organization of surfaces, based on properties of dynamic occulsion. The work on determining spatial relationships is being complemented by an investigation of techniques for estimating optical flow in situations where flow is likely to vary significantly due to boundaries or changes in surface orientation.

New York University: Robert Hummel: *Design Methods for Cooperative Labeling Processes (Computer Research)*; (DCR-8403300 A01 & A02); $54,300; 12 mos.

Labeling problems occur when a collection of distinguishable objects must be assigned individual interpretations in a consistent manner compatible with initial estimates and evidence. Examples of this very general class of problems occur in computer vision, speech recognition, cryptography, and numerous other domains of information processing. Relaxation labeling methods were introduced in the mid 1970's in order to provide a framework for studying labeling problems. Generally speaking, relaxation labeling is a process which uses local compatibility information to iteratively update an initial assignment of labels to objects so as to change either the assignment or weights representing degrees of assignments, to achieve an improved and less ambiguous labeling.

Professor Hummel and Professor Steven Zucker (McGill University) have recently developed a model characterizing the goal and the parameters in the relaxation labeling process and a convergence theorem guaranteeing reasonable behavior of the iterative procedure.

This research consists of two tasks:

1. Theoretical work involves extensions to the theory, improved convergence results, a comparison with other iterative learning and labeling network schemes, and the development of approximations to exact analytical formulas and algorithms by more "biologically plausible" methods, and
2. Empirical investigations center on applications involving edge interpretation labeling, pixel classification for effective image segmentation and analysis, and matching 3D image data structures.

New York University: David Lowe: *Model Based Recognition of Three-Dimensional Objects (Computer Research)*; (DCR-8502009); $39,400; 12 mos..

The focus of this research project is on the design and implementation of a model-based vision system that will recognize three-dimensional objects from unknown viewpoints in single gray-scale images. The capabilities of such a system are of major importance in robotics and other applications, and the methods used are also of interest to those studying the psychology of human vision. An important component of this research is the incorporation of a method for determining the exact values of viewpoint and model parameters from a few hypothesized matches between image features and object features. This allows object features to be "projected" onto the exact locations of correspondinng image features, allowing for highly robust verification of interpretations even in the presence of occlusion and missing features. Recent research results on perceptual organization and evidential reasoning are being incorporated to limit the amount of search required. The functioning system provides a framework for the incorporation of new research from many different areas of computer vision.

New York University: Jacob T. Schwartz and Herbert T. Bernstein: *Symposium on Factory Automation and Robotics, New York, New York, September 9-11, 1985*; (DCR-8511302); $10,625; 12 mos; (Joint support with the Automation and Systems Integration Program, the Computer Engineering Program, and the Manufacturing Systems Program - Total Grant $42,500).

This grant supported a symposium on Factory Automation and Robotics at New York University. The invited spokesmen represented research interests connected with parts handling, computer vision, assembly, locomotion, and robotic software. The symposium

was conducted over a three-day period and involved 150 participants. The aim of the symposium was two-fold: familiarize the theoretical and experimental computer science communities with robotics as an important area for future computer science research and provide a forum in which industrial and academic engineers and scientists in robotics interact with scientists in computer and related sciences. Problems posed by industrial applications provided fertilization across disciplinary boundaries. Applied mathematicians, engineers, and computer scientists were brought into contact with the rapid advances in robotics and factory automation. A set of videotapes of the proceedings, suitable for subsequent viewing are being produced.

University of Rochester; Dana H. Ballard; *Parameter Networks: A Connectionist Theory of Vision (Computer Research)*; (DCR-8405720 A01); $45,103; 12 mos; (Joint support with the Bioengineering & Research to Aid the Handicapped Program - Total Grant $90,206).

This project is developing a massively parallel, connectionist vision model. This model consists of upwards of 106 processors. Each processor stands for an interval in some parameter space. Physical constraints are represented by appropriate interconnections between processors. Each processor has a small amount of state that is updated iteratively as part of a relaxation process. Visual gestalts (or segmentations) correspond to fixed points in processor state space. Previous work has shown that this model can analyze complex motion and shape stimuli in only a few parallel iterations, and that this method of handling constraints, which is based on the Hough transform, is insensitive to noise and occulsion.

This research has three principal objectives:

1. Shape recognition. The objective is to extend the view transform mapping approach to shape recognition, first to connectionist models of 3D polyhedral objects, and next to connectionist models of curved objects.
2. Motion. The objective is to build an integrated connectionist motion machine that starts from raw visual data and extracts egomotion parameters.
3. Extrapersonal space. The objective is to use multiple intrinsic image cues to build an egocentric model of immediate spatial surroundings.

Results of this research are applicable to a number of application areas that could potentially help the handicapped.

Three of these areas are:

1. Visual Prostheses. This work in shape recognition and motion would lead to devices for relating this information to blind individuals via non-visual channels.
2. Visual Brain Dysfunction. A detailed characterization of the mechanisms for storing and manipulating visual information would lead to refined assessments and possibly new treatments for stroke-impaired patients.
3. Visual Learning. Knowledge about human shape representation would be invaluable in the study of dyslexia as well as the intellectually handicapped and could lead to new learning methods.

Carnegie-Mellon University; Steven A. Shafer and Takeo Kanade; *Optical Modeling in Image Understanding: Color, Gloss and Shadows (Computer Research)*; (DCR-8419990); $110,353; 12 mos.

Current methods for analyzing image intensity and shadows depend on unrealistic assumptions of ideal lighting and measurement conditions. Several analysis ideas that depend on weaker, more realistic assumptions, are being explored. Actual implementation of an environment meeting these assumptions is being developed. In this way, the theories are being tested on real images rather than on synthetic images as have traditionally been used to test quantitative theories.

The three areas being explored are color, gloss, and shadows. Color and gloss are related because highlights have a different color from the diffuse reflection from an object. Using color analysis, one can therefore separate the highlights from the diffuse reflection in an image. This produces intrinsic images useful for later stereo matching or reflectance-map analysis. Shadow analysis is being implemented under both point and extended light sources. In all these efforts, the theoretical ideas exist to some degree, and this work consists of refining these theories and testing them on real images.

This work requires an imaging environment providing a point light source and simple extended light sources, control of stray light in the room, and a carefully calibrated camera.

University of Pennsylvania; Ruzena K. Bajcsy; *Tactile Information Processing (Computer Research)*; (DCR-8410771); $46,350; 12 mos; (Joint support with the Bioengineering and Research to Aid the Handicapped Program - Total Grant $92,700).

Tactile information processing is being investigated in the context of a vision-touch integrated system. Two areas being studied are:

1. The use of touch for recognition of holes, cavities, and the surface characteristics, and
2. Tactile information processing needed for grasping.

This research has potential applications in designing computer systems to aid handicapped people so they can function normally at home, in a laboratory, or in an office setting.

University of Utah; Bir Bhanu and Thomas C. Henderson; *Computer Aided Geometric Design Based Computer Vision (Computer Research)*; (DCR-8506393); $73,451; 12 mos.

The overall objective of this project is to explore the use of Computer Aided Geometric Design (CAGD) representations and models as a basis for the visual recognition of objects for robotic applications. Specific goals are to:

1. Develop techniques and algorithms which allow the interactive and automatic generation of computer representations and geometric models of complicated realizable 3-D objects in a systematic manner. These representations and models are being obtained using available CAGD techniques and data acquired from various range finding techniques.
2. Develop fast parallel algorithms, reasoning techniques and recognition strategies which are natural to the representations and the models developed in (1) for matching 3-D models with the 3-D scene description derived from range data so as to be able to recognize, identify and locate 3-D objects including partially occluded objects.

University of Washington; Steven L. Tanimoto; *Interactive Parallel Image Processing (Computer Research)*; (DCR-8310410 A01); $64,189; 12 mos.

Parallel computers for image processing have been entering the commercial sector gradually over the last several years. The range of applications of these systems is still relatively restricted because there remains a gap in the understanding of the relationship between the operations that these machines perform (cellular logic, image filtering, etc.), and the algorithms needed to perform useful image analysis. This research employs a progression of models (for parallel image processing) to study the interactions of cellular logic operations with global image operations. The major goal is to improve the understanding of cellular logic, both conventional and pyramidal, and its interaction with overall image features.

Three tasks are being carried out to accomplish this goal:

1. Development of a capability for representing cellular logic operations in a manner which facilitates human comprehension at both theoretical and empirical levels (the representation is to be supported by an experimental software system),
2. A study of hierarchical methods in image analysis, and
3. The development of specific algorithms in the domains of industrial automation and document processing for simple, well-defined tasks.

Natural Language and Signal Understanding

University of Arkansas - Little Rock; Sally Y. Sedelow and Walter A. Sedelow; *Understanding Knowledge Representations in Dictionaries, Thesauri and Free Text (Computer Research)*; (DCR-8421195); $58,000; 24 mos.

This project pursues two parallel courses of research into the lexical and semantic structure of dictionaries, thesauri and unrestricted text with the goal of building a rich and useful integrated data base of knowledge representations for a substantial and comprehensive subset of the English language.

Analysis of dictionaries (with emphasis on a digital version of Websters's Seventh Collegiate) and text concentrates on the lexical properties of language structures that function as semantic units, utilizing a model called a collocational grammar. Lexical collocations are being identified in genre-specific language subdomains in order to account for them by means of collocational lexical-semantic relations. The same is done for the dictionary to account for knowledge representations in more general language usage.

Analysis of thesauri (with emphasis on a digital version of Roget's International Thesaurus) begins with the development of a graph or network model of the semantic structure implicit in the thesaurus' low-level clustering, using an existing mathematical model of abstract thesauri. Strength-of-semantic-association and cost-of-travel measures are being developed and used to construct a metric (if possible) spatial model of English semantic organization.

Comparison and integration of the knowledge representations derived from dictionaries, thesauri and text are used to facilitate the development of a relatively comprehensive knowledge base, which is being tested by the creation of intelligent systems for text generation and understanding.

SRI International; Douglas Appelt; *Natural Language Utterance Planning (Computer Research and Information Science)*; (DCR-8407238); $88,804; 12 mos; (Joint support with the Information Technology Program - Total Grant $111,504).

This research project is concerned with the problems of formalizing and planning utterances in natural language processing. Formal methods developed for utterance planning are being extended to provide a more

detailed account of different types of referring actions. This investigation consists of:

1. Developing the underlying logical theory necessary to support reasoning about referring actions;
2. Using the logical theory to formalize referring actions as operators in a planning system; and
3. Augmenting a grammar of English so that an utterance planner can produce sentences that realize different types of referring actions.

Illinois Institute of Technology; Martha W. Evens; *Automatic Construction of Lexicons with Relational Primitives (Information Science and Computer Research)*; (IST-8510069); $26,235; 24 mos; (Joint support with the Division of Information Science and Technology - Total Grant $80,240).

Natural language processing systems need access to much larger lexicons, or dictionaries, than those available today. Such systems also require information about the lexical and semantic relations among words. Such information is not usually represented explicitly in commercial dictionaries. This research project is investigating problems in the theory of lexical-semantic relations. Procedures for automatically constructuring lexicons for substantial sublanguages are being developed; these are based both upon automatic extraction of information from machine readable dictionaries and upon interaction with human experts.

The significance of this research lies in its contribution to the design of lexicons to support computer systems with natural language processing capabilities.

Colgate University; Allen B. Tucker; *Knowledge-Based Multilingual Machine Translation (Computer Research and Information Science)*; (DCR-8407114 A01); $35,180; 12 mos; (Joint support with the Division of Information Science and Technology - Total Grant $51,180).

The early extensive effort in machine translation (MT) which started with Weaver's memorandum (1949) did not produce adequate results, although much was contributed to the development of artificial intelligence and computational linguistics. The work on MT continued with much less funding and with emphasis mainly on operational, non-theoretical issues. The developments gathered momentum with the spectacular improvements in the computational support over the past 20 years, so that at present several MT systems are in operation. Recently, several MT projects have been initiated in the European Economic Community, Japan, Israel and Soviet Union. Significant development of the science and technology of artificial intelligence and knowledge-based expert systems has led to the emergence of the proposition of a knowledge based MT system which has a substantial real world knowledge base in addition to knowledge bases of the source and target languages.

A knowledge-based model for multilingual MT is being developed and tested. The model combines linguistic, morphological, and syntactic parsing technology in current MT systems with "world knowledge" and "short-term memory" functions of recent artificial intelligence research.

Three major goals of the research are to develop and test:

1. A universal intermediate language (IL) for representing the syntactic and semantic information in a text and the knowledge base itself;
2. A knowledge base for a specific subject domain, written in IL, which is well adapted to multilingual MT, and contains linguistic knowledge as well as world knowledge; and
3. An objective measure of the quality of a MT system.

An English corpus of 10,000 words in machine readable form is being used as a basis for this experiment.

New York University; Ralph Grishman (in collaboration with Lynette Hirshman at the Burroughs Corporation); *Industry/University Cooperative Research: Acquisition and Use of Semantic Information for Natural Language Processing (Computer Research)*; (DCR-8501843); $43,250; 12 mos; (Joint support with the Division of Industrial Science and Technological Innovation - Total Grant $86,500).

The acquisition and use of semantic information is a continuation of an ongoing research effort focused on robust parsing of natural language input in limited domains. This research addresses the problems of portability, reliability, and habitability. It outlines the development of a set of interactive tools which provide, to both the linguist and the end-user, an accurate set of domain-specific co-occurrence patterns (subject-verb-object and host-modifier combinations allowed in the domain). This set of tools will partly automate the creation of the Domain Information Schema (DIS), namely the domain-specific information required for processing a given domain. This leads to increased portability of the processing system, since specification of the domain-specific information is a very labor-intensive process. The same set of tools, provided with a different interface for the end-user, are used to support interactive text entry. The tools provide detailed feedback on where the system fails to process the text, due to missing or incompatible semantic co-occurrence patterns. This increases the habitability of the system by making it easier for the user to compose text which the system understands. As the user adds the missing patterns to the system, the robustness of the system should improve as well.

This research is proceeding at two sites, NYU (interactive text entry) and Burroughs Corporation (creation of the DIS). The project consists of three stages: tool-building and integration; initial test on one or two

domains and refinement of tools; and final test of new domain(s), and evaluation of tools.

Burroughs Corporation; Lynette Hirshman (in collaboration with Ralph Grishman at New York University; *Industry/University Cooperative Research: Acquisition and Use of Semantic Information for Natural Language Processing (Computer Research)*; (DCR-8502205); $34,606; 12 mos; (Joint support with the Division of Industrial Science and Technological Innovation - Total Grant $69,211).

The acquisition and use of semantic information is a continuation of an ongoing research effort focused on robust parsing of natural language input in limited domains. This research addresses the problems of portability, reliability, and habitability. It outlines the development of a set of interactive tools which provide, to both the linguist and the end-user, an accurate set of domain-specific co-occurrence patterns (subject-verb-object and host-modifier combinations allowed in the domain). This set of tools will partly automate the creation of the Domain Information Schema (DIS), namely the domain-specific information required for processing a given domain. This leads to increased portability of the processing system, since specification of the domain-specific information is a very labor-intensive process. The same set of tools, provided with a different interface for the end-user, are used to support interactive text entry. The tools provide detailed feedback on where the system fails to process the text, due to missing or incompatible semantic co-occurrence patterns. This increases the habitability of the system by making it easier for the user to compose text which the system understands. As the user adds the missing patterns to the system, the robustness of the system should improve as well.

This research is proceeding at two sites, NYU (interactive text entry) and Burroughs Corporation (creation of the DIS). The project consists of three stages: tool-building and integration; initial test on one or two domains and refinement of tools; and final test of new domain(s), and evaluation of tools.

University of Pennsylvania; Aravind K. Joshi; *Research In Natural Language Processing (Computer Research)*; (DCR-8410413); $153,000; 12 mos.

Natural language processing is an important area of current research in intelligent systems. A major part of this research deals with several aspects of cooperative behavior in the context of interaction with databases and expert systems. In particular, cooperative responses related to object related misconceptions, interaction with dynamic database systems, and preventing false inferences are being investigated. Other parts deal with certain aspects of temporal reasoning, certain discourse phenomena, and certain constrained formal systems.

University of Pennsylvania; Insup Lee and Ruzena Bajcsy; *A Distributed Testbed for Active Sensory Processing (Computer Research)*; (DCR-8501482); $66,208; 24 mos.

Although distributed processing appears to have many benefits in terms of performance, resource sharing, and adaptability, the scientific community has relatively little experience in the implementation and the actual use of distributed systems for real-time applications. This research is attempting to design and implement an experimental distributed system composed of a heterogeneous mix of computers. This system provides a coherent distributed environment for the investigation of the architectures and algorithms for multi-sensory (vision and tactile) analysis and manipulation of three-dimensional objects. The objectives of this research are:

1. To gain a better understanding of the organization and structure of a heterogeneous distributed system for real-time applications; in particular, for a multi-sensory system;
2. To evaluate the practicality of known strategies for the dynamic reconfiguration of a distributed program through real (not simulation) experiments, and to propose better strategies that can be used by distributed system designers to improve the system resource utilization and response time of their systems; and
3. To design and implement a distributed programming system that facilitates the construction and execution of a distributed program whose components are written in different programming languages like C, FORTRAN, LISP and Prolog.

Automatic Theorem Proving, Concept Learning, and Inference

Stanford University; John McCarthy; *Artificial Intelligence (Computer Research)*; (DCR-8414393); $121,619; 12 mos.

The main long term goal of this research has been finding ways of formalizing common sense knowledge and reasoning ability so that a computer program can have the common sense knowledge and the common sense ability to achieve the goals of the system. Expressing common sense knowledge in languages of mathematical logic and using controlled deduction as a way of deciding what to do was first proposed by

Professor McCarthy. A major advance in the use of logical languages for artificial intelligence was the notion of formalized non-monotonic reasoning which was proposed also by Professor McCarthy. Non-monotonic reasoning and the circumscription approach to formalizing it have wide applications not only in artificial intelligence but potentially also in databases, programming languages and even in analytic philosophy. This project continues to explore these new possibilities.

Yale University; Dana C. Angluin; *Algorithms for Inductive Inference (Computer Research)*; (DCR-8404226); $59,196; 12 mos.

Inductive inference is the process of guessing general rules from examples. Previous work includes efficient inductive inference methods for specific domains, methods for diagnosing and correcting errors in Prolog programs, specifying transformations in a text editor by means of examples (editing by example), and probabilistic inductive inference. This research focuses on the efficiency of methods for inductive inference in concrete domains, with emphasis on potential applications, probabilistic methods, achieving generality by means of explicit representation of domain-specific information, and expanding the basic paradigm of the field to be able to model a wider range of phenomena in this area.

Georgia State University; Arthur J. Nevins: *Hierarchical Learning of Multiple Concepts (Computer Research)*; (DCR-8408389); $70,830; 24 mos.

One of the important problems in current research in artificial intelligence is knowledge acquisition. In expert systems research, for example, where recent advances have made it possible for computer systems to emulate human expertise and problem solving skills in a particular knowledge-based field, there is a critical need to develop methods to acquire expertise or knowledge automatically and efficiently. This is a difficult and challenging problem facing the AI community today. This research project is exploring effective ways to enable a computer to acquire the sophisticated problem solving techniques of an experienced mathematician. The goal is to combine deductive and inductive capabilities in a single problem solving system.

Northwestern University; Lawrence J. Henschen; *Logic and Databases (Computer Research)*; (DCR-8306637 A02); $67,326; 12 mos.

The purpose of this research is to enhance the information retrieval capabilities of data bases by incorporating first-order logic into the data base. The relational data model is being used in which the relations themselves already are of the form of elementary logical statements. With this model, it is only necessary to supply the appropriate logical mechanisms for integrating more general logical formulas into the data storage and retrieval processes. It has been shown such a model can be used when the logical formulas represent definitions of new relations, even when such definitions are recursive. This allows the possiblity of greatly reducing storage because such defined relations need not be stored; rather, their information content can easily and efficiently be generated when needed. This can be accomplished at data base creation time, a major step forward. In this study other types of formulas (e.g., constraints) are being examined to determine whether they can be "compiled" at data base creation time. The research is also looking into how functions, which normally lead to infinite deduction paths, can be handled in the specialized domain of data bases.

University of Illinois - Urbana; R. S. Michalski; *Studies in Computer Inductive Learning and Plausible Inference*; (DCR-8406801 A01); $125,000; 12 mos.

This research is concerned with computer inductive learning and formal models of human plausible inference. It encompasses basic theoretical studies as well as the development of algorithms and their computer implementations and experimental applications to selected practical problems. Major topics of the research include:

1. Automated generation of derived descriptors (variables, relations, functions) which are most relevant for constructing inductive assertions;
2. The development of rules that govern the generalization processes and efficient strategies for their application in inductive inference;
3. The study and implementation of techniques for constructing conceptual hierarchies of objects;
4. The development of experimental multi-purpose high performance inductive learning systems which incorporate the methodology for generating derived descriptors; and
5. The application of inductive inference to the design of computer architectures and compiler structures.

Brandeis University; Jacques Cohen; *Compiler Prototyping Using Logic Programming (Computer Research)*; (DCR-8500881); $24,000; 12 mos; (Joint support with the Software Engineering Program - Total Grant $66,982).

This research is applying techniques developed using Prolog to current topics in compiler development, including parallel compilation and code generation. Equally important is the investigation of the development of efficient compilers and interpreters for extensions of Prolog itself. The goal of the research is to produce Prolog programs for compiler applications that will execute as efficiently as typical C programs.

Harvard University; Leslie G. Valiant; *Parallel Computation (Computer Research)*; (DCR-8542868); $147,000; 12 mos.

The complexity limitation of computational devices to learn concepts has been investigated by Professor Valiant. Although the ultimate aim is to consider parallel devices it was discovered that no satisfactory precise specifications of learning were available even on sequential devices. The question of whether the class of conjunctive concepts might be learned in feasibly, few computational steps could not even be asked. Substantial progress has been made both in discovering plausible specifications of the concept learning process and in discovering non-trivial classes of concepts that can be learned efficiently.

Towson State University; John Grant and Jack Minker; *Research in Database Foundations and Logic Programming at An Undergraduate Institution (Computer Research)*; (DCR-8412662); $79,150; 36 mos.

One of the important areas of research in artificial intelligence (AI) is knowledge-based systems. These systems deal with large and often distributed databases with complex organizations capable of storing and manipulating highly specialized knowledge in a specific domain. Recent work in AI has shown that there is a strong link between database reseach and research in knowledge-based systems. This link is made possible because of the potential application of mathematical logic and logic programming to databases. In this project, Professors J. Grant and J. Minker are investigating several important areas in this approach - query optimization, transition constraints, generative vs deductive axioms, null values and indefinite databases.

University of Maryland - College Park; Jack Minker; *Artificial Intelligence, Parallel Logic Programming, and Deductive Databases (Computer Research)*; (DCR-8305992 A02); $54,136; 12 mos; (Joint support with the Division of Information Science and Technology - Total Grant $108,271).

The major thrust of this research is the application of automated theorem proving and mathematical logic to database investigations.

Three types are being treated:

1. Theorem proving by linear resolution with unrestricted selection function using trees,
2. First-order logic and databases, and
3. Database logic.

Database logic is a recently developed formalism whereby existing database systems can be viewed in terms of logic. It has the prospect of providing a unifying theoretical framework for database systems.

Rutgers University; Tomasz Imielinski; *Processing Incomplete Knowledge-A Database Approach (Computer Research)*; (DCR-8504140); $54,607; 12 mos.

Representing, manipulating and processing incomplete information or knowledge are increasingly important to database management and knowledge engineering. The reason is that often knowledge is incomplete and information is under-specified. The issue which is particularly challenging is the problem of query processing in the presence of incomplete data.

This problem has two aspects: specification and implementation. The first is to provide the proper language together with the semantics to express queries and has been a topic of considerable interest in artificial intelligence (AI). Powerful specification languages have been proposed; however, the problem of query processing is left open. The second has been the focus of attention in the field of database management systems (DBMS). In contrast to other AI research, only very limited types of techniques for dealing with information incompleteness have been considered. This research is combining the specification-oriented AI approach with the abstract implementation- oriented DBMS approach so that efficient methods for query processing in large databases of incomplete information can be obtained.

Carnegie Mellon University; Peter B. Andrews; *Automated Theorem Proving in Type Theory (Computer Research)*; (DCR-8402532 A01); $88,318; 12 mos.

This project is directed toward enabling computers to construct and to check proofs of theorems of mathematics and other disciplines formalized in type theory (higher order logic) or first order logic, and to assist humans engaged in these tasks.

Previous research has shown that one can search for a proof of a theorem of higher order logic by searching for an expansion tree proof (ET-proof) of the theorem, and then transforming this into a proof in natural deduction style. This research focuses on developing and implementing a procedure for finding ET-proofs, and on related theoretical questions.

An existing computerized theorem proving system called TPS is being extended to facilitate generating and manipulating examples useful in this research, and to translate back and forth between ET-proofs and natural deduction proofs. TPS is also being enhanced as a practical and convenient tool for constructing and checking proofs interactively, semi-automatically, or automatically.

Carnegie Mellon University; Elaine Kant and Allen Newell; *Algorithm Design and Discovery (Computer Research and Information Science)*; (DCR-8412139); $70,000; 12 mos; (Joint support with the Information Technology Program - Total Grant $140,000).

One of the goals of artificial intelligence (AI) is to design and synthesize algorithms and programs. Professors E. Kant and A. Newell are implementing an AI

program that designs algorithms and programs from specifications of the inputs and desired outputs as well as gaining a better understanding of the human algorithm-design process. This research is an extension of the ongoing "Designer" project at Carnegie-Mellon University. The following areas are being investigated:

1. Automatic design of computational geometry algorithms,
2. Application of the "divide annd conquer" method to design algorithms in an algebraic domain,
3. Algorithm analysis as a means of guiding design, and
4. Transformation of algorithms to code automatically.

St Joseph's University; Ranan B. Banerji: *Knowledge Based Learning and Problem Solving Heuristics*; (DCR-8504223); $71,400; 24 mos.

This research project is concerned with knowledge-based learning. The goal is to develop both theory and experiments into a learning algorithm which incorporates the strengths of many existing systems. One of the systems developed by the PI and initiated by Pennypacker has been based on the insight that the efficiency of learning is influenced by the knowledge of the learner. Thus, there is a need for incorporating into a learning system a method by which the knowledge, as acquired, can be used in increasing the efficiency of the learning. Most of the learning systems have their learning algorithm frozen at the point of design and do not take advantage of newly acquired knowledge. Contrariwise, Pennypacker's early work, even though it was based on propositional logic, had partially succeeded in using acquired knowledge to guide new learning. More recently, his method was enhanced by Cohen and Sammut to incorporate a predicate calculus with function symbols and to be capable of learning the descriptions of relations in addition to those of sets. Many of the descriptions learned were recursive, since the learning algorithm was free to use previously learned incomplete descriptions of a concept to be incorporated into more complete descriptions. Except for the lack of knowledge dependence, the alternative approaches to learning developed by Mitchell and by Shapiro have proved to be insightful by supplying formal theoretical bases for the study of learning processes. These approaches are being analysed and synthesized in a unified framework.

Brown University; Ulf Grenander; *Inference in Markov Processes (Mathematical Sciences and Computer Research)*; (DMS-8541864); $8,500; 12 mos; (Joint support with the Statistics and Probability Program - Total Grant $22,700).

This research deals with the mathematical foundation for models of pattern recognition and image restoration, an area where Professor Grenander has long done notable work and has attracted considerable attention. Substantial progress has been achieved during the past year, including formulation of estimation methods for pattern recognition models, generation of random geometries, and the development of statistical software. This project continues research along these lines.

University of Texas - Austin; Woodrow W. Bledsoe; *Automatic Theorem Proving and Applications (Computer Research)*; (DCR-8313499 A01); $144,273; 12 mos.

This research is a continuation of present work in automatic theorem proving and its applications to program verification and mathematics. Earlier work has included the development of computer programs which: prove theorems as part of a program verification system; prove theorems from intermediate analysis with techniques for instantiating set variables and for handling general inequalities. Other work has been on complete sets of reducers and agenda driven provers. This investigation continues the work, especially in program verification and complete sets of reducers, and continues the effort to build a more powerful general theorem prover.

University of Texas - Austin; Robert S. Boyer and J. Strother Moore; *Mechanical Proving in Geometries (Computer Research)*; (DCR-8503498); $82,600; 24 mos..

The focus of this research project is on mechanical proving of theorems in various geometrics using the algorithm discovered by Wu Wen-tsun, the leading mathematician in China. The geometrics are: Euclidean plane geometry and solid geometry, projective plane geometry, Bolyai-Lobachevsky's hyperbolic geometry, Riemann's elliptic geometry, Mobius geometry (inversion geometry), Lagerre's geometry (equilong geometry) and Kanser's conformal geometry. The study consists of three subgoals:

1. Implementing provers for various geometrics to provide known theorems;
2. Using several techniques to discover and prove new theorems; and
3. Developing systems of axioms of some geometrics (e.g. Bolyai-Lobachevsky's geometry) suitable for Wu's methods, thus establishing the connection between Wu's method and the method used in the main stream of automated theorem proving: formal decisions.

Knowledge Representation and Problem Solving

Stanford University; Edward A. Feigenbaum and Charles Yanofsky; *MOLGEN-Applications of Artificial Intelligence to Molecular Biology Research in Theory Formation, Testing, and Modification (Computer Research)*; (DCR-8310236 A01); $131,621; 12 mos.

The MOLGEN project is exploring the cognitive process of scientific theory formation with the aim of building automated reasoning tools that can assist a scientist in this complex endeavor. The purpose of much of experimental science is to find explanations for natural behavior. These explanations take the form of theories: Detailed models that can explain observed behavior and predict as yet unobserved behavior. The task of constructing, modifying, extending, and discarding, occasionally, scientific theories is the object of this research. The research consists of several major phases. The first, recently completed, was a detailed study of the reasoning processes to discover the regulatory processes operating in the trp operon system of E. coli. The second phase is to build a knowledge base that adequately represents the currently accepted theory of the biological system; adequacy here refers to the ability to explain laboratory observations and predict logical consequences that result from the theory. The third phase involves creating the reasoning mechanisms that permit theory modification, extension and complete re-creation, occasionally within the context of the Yanofsky trp operon system. Finally, results are being generalized and verified by applying the theory formation system to a different biological system.

Stanford University, Edward A. Feigenbaum; *The Mechanization of Formal Reasoning (Computer Research)*; (DCR-8303142 A02); $102,223; 12 mos.

This project continues work on the mechanization of formal reasoning for artificial intelligence systems. Based on past research, it focuses on three new and extended components:

1. Continued theoretical research in FOL, a system capable of three different levels of reasoning: Reasoning at the object level, meta-level and self-referring;
2. Experimental development of an expert system using the FOL framework; and
3. Exploration of the role of formal reasoning in the design of high-level architecture for artificial intelligence.

Stanford University; Oleg Jardetzky and Bruce B. Buchanan; *Interpretation of NMR Data From Proteins Using Artificial Intelligence Methods*; (DMB-8402348); $50,000; 24 mos; (Joint support with the Biophysics Program - Total Grant $100,000).

The solution structure and dynamics of proteins remains an unsolved and elusive problem. Though good progress has been made in understanding and developing methods for the study of secondary structure, hydrogen exchange and fast chain motions, knowledge of the tertiary structure and of slow motions is in a much less satisfactory state.

Dr. Jardetzky has been examining the possibilities of deriving information on the tertiary structure of proteins and on the slow motions within them that are important to function from NMR data. Along with Dr. Buchanan, he is currently testing new types of experiments to bear on these problems and to formulate an artificial intelligence program to correlate existing information to find the correct solution(s) by event-driven reasoning.

University of California - Berkeley; Lotfi A. Zadeh; *Fuzzy Logic as a Basis for Commonsense Reasoning and Inference in Expert Systems (Computer Research)*; (DCR-8513139); $93,963; 12 mos.

The principal aim of this research is to explore the suitability of fuzzy logic as a framework for commonsense knowledge representation, commonsense reasoning and the management of uncertainty in expert systems.

Fuzzy logic differs from the traditional logic systems in three basic respects. It allows the use of fuzzy predicates and linguistic variables exemplified by small, large, tall, fast, red, not very tall, quite tall and extremely tall; fuzzy truth values exemplified by not very true, somewhat true and almost true; and fuzzy quantifiers exemplified by most, almost all, many, few, several, often and usually.

An important class of problems which arise in the context of commonsense reasoning is that in which the facts and rules in the knowledge base are of the form: usually X is F and if X is F then usually Y is G, where F and G are fuzzy predicates, X and Y are constrained variables, and usually plays the role of a conditioned or unconditioned fuzzy quantifier. One of the basic aims of this research is that of developing a systematic method for dealing with propositions of this type both in the context of commonsense reasoning and in the management of uncertainty in expert systems.

University of California - Los Angeles; Judea Pearl; *Studies in Heuristics (Computer Research)*; (DCR-8501234); $77,306; 12 mos.

The main objective of the research project is to continue studies leading towards a better understanding of the principles which govern the utilization of heuristic models in problem solving. To be quantified are the relationships between the precision of the heuristics used, the quality of the solutions they produce and the complexity of search which they guide. Means by which heuristics models can be generated mechanically and schemes which would allow the search to be conducted in parallel are also being investigated.

University of California - Los Angeles; Judea Pearl and Moshe Ben-Bassat; *Toward a Computational Model of Evidential Reasoning*; (DCR-8313875 A01); $89,672; 12 mos.

This research program is directed at the development of computational models for evidential reasoning; that is, inference mechanisms by which the evidence provided by a set of findings is analyzed in order to gain better understanding of a given situation or phenomenon. Problem solving tasks of this nature include, for instance, medical diagnosis, weather forcasting, corporate assessment, political crisis assessment and battlefield reading. These problems are referred to as situation assessment tasks.

The approach is based on the decomposition of the situation assessment task into evidence integration, hypotheses generation and evaluation, goal(s) setting, information sources evaluation and selection, and sorting of evidence by goals and hypotheses. Inference nets are used for knowledge representation in which the nodes represent observable and inferred events, while a link between nodes Ei and Ej indicates evidential relevancy between the two links. Each link is assigned value(s) that represent the degree of significance for inferring Ei from Ej.

Yale University; Paul Hudak; *DAPS: Systems Support For AI (Computer Research)*; (DCR-8403304); $20,000; 12 mos.

Parallel computers used in scientific computation are successful because of their ability to take advantage of the regularity and predictability of algorithms in the field. However, Artificial Intelligence (AI) programs are characterized by their irregularity and unpredictability, for which existing machines are poorly suited. Professor Hudak is investigating fundamental similarities in the dynamic behaviors of many AI programs so that computer systems can be designed accordingly. In phase one, an empirical study of a representative set of AI programs is being conducted to determine the nature of dynamic factors. In the second phase, heuristic strategies are being embedded into a class of machines called DAPS that is tailored to deal with AI program behavior.

Yale University; Drew V. McDermott; *Spatial Reasoning for Problem Solving (Computer Research)*; (DCR-8407077 A01); $51,501; 12 mos; (Joint support with the Bioengineering and Research to Aid the Handicapped Program- Total Grant $103,001).

The goal of this project is to investigate approximate shape representations to support reasoning about changing systems. Previous work on "fuzzy maps" is being generalized so that the fuzzy numbers are attached to shape representations at different levels of resolution. This research is attempting to formalize what it means for an object to have multiple shape representations. Surface representations are being treated as primary, and exactness is being characterized by two parameters: the "gap size," which bounds the size of holes neglected in abstracting this surface; and the "grain size," which bounds the discrepancies between the idealized surface and the actual surface. Such approximate shape representations can be used to support reasoning about motion, shape change, and fluid flow. They provide a convenient way of abstracting "faces" from arbitrary shapes. These faces can then be the focus of reasoning about degrees of freedom of motion, fluid containment, and so forth. This project is producing results useful in robotics research or robotic applications such as computer aids for the handicapped.

If handicapped people are to profit from robotics, they will need mobile robots that have the same ability to contemplate and manipulate representations of a changing, incompletely known world. Such robots could be used in two ways: as autonomous servants, and as direct symbionts of a disabled person. In the first mode, the robot would wander about a house or workplace detached from its user. It would respond to commands to go and retrieve objects, or perform simple tasks like unlocking a door to let someone in. In the second mode, the robot would become part of its user, controlling prosthetic legs and hands. This research is directed toward developing the spatial and temporal reasoning that will be needed by such mobile robots.

Illinois Institute of Technology; Shi-Kuo Chang; *An Intelligent System for Crisis Management (Computer Research)*; (DCR-8306282 A01); $71,565; 12 mos..

This project examines the methodology for designing intelligent systems for the management of imprecise data, database alerters and message filters, with applications to crisis management. The focus is on the study of intelligent filters for message management, of office procedure analysis and their experimentation. Linguistic filters used in the experimental message management system are being extended to intelligent filters by the incorporation of alerter rules. Office procedure models are being investigated and extended to timed k-bounded Petri-net analysis.

University of Maryland - College Park; Laveen N. Kanal; *Parallel Problem Solving and Applications in Artificial Intelligence*; (DCR-8504011); $71,896; 12 mos.

This research focuses on selected parallel processing approaches and algorithms with applications in artificial intelligence (AI). The research has evolved from the recent development of unifying approaches and parallel algorithms in AI and Pattern Analysis. The aim is to develop hybrid versions of certain search algorithms to obtain efficient, parallel algorithms for problem solving tasks. These techniques are being extended to more general task graphs in multiprocessing. Given the task graph of a problem for a specific problem size, the automatic generation of the graph for a larger problem size is very useful in the design of parallel algorithms. Another goal of the project is to study how this hybrid representation can be used to infer many important properties of classes of solution strategies. Finally, the results obtained are being applied in the analysis of signal ensembles and design of hierarchical decision making structures.

University of Massachusetts - Amherst; Victor R. Lesser; *Coordination in Cooperative Distributed Problem Solving (Computer Research)*; (DCR-8300239 A02); $64,348; 12 mos.; (Joint support with the Computer Systems Design Program and the Special Projects Program - Total Grant $128,696).

A key problem in cooperative distributed problem solving is developing network coordination policies that provide sufficient global coherence for effective cooperation. This problem is difficult because limited internode communication precludes the use of a global "controller" node. Instead, each node must be able to direct its own activities in concert with other nodes, based on incomplete, inaccurate and inconsistent information. This requires a node to make sophisticated local decisions that balance its own perceptions of appropriate problem solving activity with activities deemed important by other nodes. Current research at the University of Massachusetts has included the development of a node architecture capable of such sophisticated local decision making. This architecture has been implemented as part of the distributed vehicle monitoring testbed: a flexible and fully instrumented research tool for the empirical evaluation of distributed network designs and coordination policies. This research builds on this node architecture and testbed to explore (through actual implementation and empirical studies) a variety of approaches to network coordination that include: organizational self-design, distributed load-balancing, negotiation among nodes, planning of internode communication strategies, and knowledge-based fault-tolerance.

University of Massachusetts - Amherst; Victor R. Lesser; *A Research Facility for Cooperative Distributed Computing (Computer Research)*; (DCR-8500332); $600,000; 12 mos; (Joint support with the Coordinated Experimental Research Program - Total Grant $1,237,941).

This project provides an experimental facility for research in cooperative distributed computing based on a tightly-coupled multiprocessor containing between 64 and 128 processing elements. A key aspect of the project is to address issues that arise out of task and problem solving decomposition in which processes cooperate to solve a single integrated problem rather than a set of independent tasks. This project coordinates and integrates large, successful research groups in the areas of intelligent systems (computer vision, robotics and distributed problem solving, adaptive learning networks, neural modeling) and distributed and parallel processing (language primitives, synchronization, concurrency control, network communications, real-time scheduling, load balancing, debugging, protection and reliability).

University of Michigan - Ann Arbor; Arthur W. Burks; *Languages and Architectures for Parallel Computing with Classifier Systems (Computer Research)*; (DCR-8305830 A02); $105,042; 12 mos; (Joint support with the Computer Systems Design Program - Total Grant $105,042).

A new class of rule-based computing systems, classifier systems, has been developed and tested. While classifier systems are similar to production systems, they are different in basic ways which facilitate the incorporation of learning algorithms in them, and which make it possible for them to be executed in a highly parallel fashion without the usual problems of concurrency and interlock. Classifier systems with learning algorithms have been tested successfully on two model control systems, a game learning program, and a hardware robot project. Progress has been made on the theoretical understanding of how such classifier systems work.

This research focuses on two interrelated areas:

1. The powers and limits of classifier systems are being studied by further computer simulations and theoretical analyses; and
2. Novel parallel architectures for the rapid execution of classifier systems are being planned and simulated.

On the theoretical side, these studies contribute to a deeper understanding of concurrency and distributed control, and on the practical side they lead ultimately to powerful parallel computers which are non-von-Neumann in both programming and architecture.

University of Michigan - Ann Arbor; Douglas R. Hofstadter; *The Copycat Project: An Experiment in Nondeterminism and Creative Analogies (Computer Research)*; (DCR-8410409); $88,820; 24 mos.

In human learning and reasoning, many analogies involving different concepts and levels of subtlety can be made. The problem of how to implement a computational model of this human ability is considered. A key part of the proposed system is its dependence on statistically emergent properties of stochastically interacting "codelets" (small pieces of ready-to-run code created by the system, and selected at random to run with probability proportional to heuristically assigned "urgencies"). Another key element is a network of linked concepts of varying levels of "semanticity", in which activation spreads and indirectly controls the urgencies of new codelets. There is pressure in the system toward maximizing the degree of "semanticity" or "intensionality" of descriptions of structures, but many such pressures must interact with one another, and compromises must be made.

University of Minnesota; James R. Slagle; *Expert Systems Questioning Procedures Based on Merit (Computer Research)*; (DCR-8512857); $79,263; 12 mos.

Expert consultant systems often perform computations on a directed graph of associated propositions. Each proposition is represented by a node. Edges connecting these nodes are associated with rules which organize the propositions into antecedent/consequent relationships. A node may be assigned a value through the edges that bind it to its antecedents. The merit strategy is employed to determine assignment sequences that result in efficient computer consultation by focusing the consultation process on the most appropriate questions. The merit system, originally defined for logical functions in the MULTIPLE program, has been extended to the MYCIN style of propagation and to the method of subjective Bayesian assignments used by PROSPECTOR. The experience has shown that merit values provide an efficient flow of control for expert consultation. This research is attempting to extend merit such that it can be practical for a large knowledge base, a large database, directed acyclic graphs, and propagations of distributions. Merit should also be able to handle a potential finite change to the value, and should be defined and used relative to any ancestor node, not just to a top node.

Rutgers University - Busch Campus; Saul Amarel and Charles Schmidt; *Exploration of Problem Reformulation and Strategy Acquisition*; (DCR-8318075 A02); $79,852; 12 mos; (Joint support with the Division of Information Science and Technology - Total Grant $106,470).

The problem solving ability of an Artificial Intelligence (AI) program is critically dependent on the nature of the symbolic formulation of the problem given to the program. Improvement in performance of the problem solving program can be made not only by improving the strategy of directing search but also by shifting the problem formulation to a more appropriate form. Certain formulations are more amenable to incremental acquisition of problem solving strategy than others. With this in mind, an extensible problem reduction method is being developed that allows incremental strategy construction. The overall objective of this research project is to study strategy acquisition and problem reformulation for planning problems.

Cornell University; Robert L. Constable; *Experiments with a Program Refinement System (Computer Research)*; (DCR-8303327 A02); $63,000; 12 mos; (Joint support with the Software Engineering Program and the Software Systems Science Program - Total Grant $189,494).

This research involves the judicious combination of the strong points of program development methodology and program verification in a logic for correct program development. This project is exploring a logic for formal top-down development called refinement logic. A system called Program Refinement Logic (PRL) has been designed and implemented. This system is being improved and expanded and considerable experimentation is being conducted.

Cornell University; John E. Hopcroft and Kuo K. Wang; *A Program of Research in Representing Physical Objects*; (DCR-8502568); $99,858; 12 mos.

This project is part of a multidisciplinary research effort in developing computer representations of physical objects and systems. While this project concentrates on the computer science and mechanical engineering aspects of the representation problem, the group is working with researchers investigating related topics.

The research is aimed at:

1. Investigating methodologies, algorithms and mathematical foundations for describing and reasoning about physical objects and systems,
2. Designing and implementing the skeleton of a solid modelling system that can support abstract modelling and surface representations of quadric or higher complexity, and
3. Designing and implementing an editing system that supports functional specification and modification of representations of physical objects, assemblies and systems.

Cornell University; John Hopcroft and Alan Demers, *A Program of Research in Robotics*; (DMC-8312096 A01); $87,291; 12 mos; (Joint support with the Computer Engineering Program - Total Grant $261,822).

This is a multidisciplinary program of research in robotics.

The research program has two components:

1. A design and implementation component to provide a robotic environment in which to validate new algorithms and ideas on programming language and environment design, and
2. A theoretical program of research aimed at developing the underlying mathematical foundations necessary for robotics.

The design and implementation efforts have centered on developing a system to allow experimentation with a number of different aspects of the robot programming environment. One of the most important of these is the problem of object representation for single objects and generic classes of objects as well. For example, the system should allow a user to describe the concept "bolt" independent of any specific example. Other examples of areas being investigated are the problems of motion and task planning algorithms, generic transformations (e.g. bolting, welding), and graphics-based user interfaces. The initial environment is directed toward automated assembly, but the system serves much more as a test-bed for new ideas in language and algorithm design than as a production robot programming environment.

New York University; Ernest Davis, *Reasoning about Shape and Function (Computer Research)*; (DCR-8402309); $50,000; 24 mos.

This research project is developing an artificial intelligence theory of reasoning about the relations between the shapes and the uses of three-dimensional rigid objects. The representation of three-dimensional shapes and motions, both quantitative and qualitative, and algorithms to manipulate these representations are being investigated.

New York University; Colm O'Dunlaing and Chee-Keng Yap; *Motion Planning Problems in Robotics: Algorithmic Issues (Computer Research)*; (DCR-8401898 A01); $90,600; 12 mos.

Motion planning is a central problem in robotics. As a purely algorithmic problem, it has recently been investigated extensively by a number of researchers.

This work is continued and extended in some new directions, notably:

1. Development of the retraction method. This powerful technique has given rise to the most efficient algorithms known in each case studied. Although its general applicability to motion planning problems seems clear, this remains to be proved.
2. Motion with bounded acceleration. Some initial work has begun by considering the problem of finding a motion with minimal maximum acceleration sufficient to avoid a given set of moving obstacles.
3. Coordination of motion. Near optimal algorithms for the cases of coordinating two or three discs have been given. Practical refinements of these algorithms and the coordination of two arms in space are being considered.
4. Experimental issues. Implementation of the retraction algorithm for a moving disk is under way. Two issues are being studied: the "robustness" of the method and practical methods to shorten the path obtained through the Voronoi diagram.

University of Rochester; James F. Allen; *Temporal World Models for Problem Solving (Computer Research)*; (DCR-8502481); $35,601; 12 mos.

Several related issues in the representation of temporally qualified knowledge and its application to problem solving in complex environments are being considered. This research project is concerned with problems that arise when only partial information is available about the world, and the planner may have to plan to interact with external events.

In particular, there are three main goals:

1. To develop a formal theory of belief, choosability, and possibility within the framework of the existing interval-based temporal logic,
2. To use this theory to guide the implementation of a prototype planner that can reason about uncertainty and plan to interact with external events; and
3. To examine the foundations of interval-based temporal logics, concentrating on producing a concise axiomatization of the theory and examining the relationships between time intervals and pseudo-time points.

University of Texas - Austin; Benjamin J. Kuipers; *Knowledge Representations for Expert Causal Models (Computer Research)*; (DCR-8512779); $84,700; 12 mos.

The goal of this research is to design and build a computer program to understand a realistic description of a medical case, construct a detailed model of the internal state of the patient and its evolution, and answer questions about the pathophysiology of the case, concentrating on the domain of salt and water balance disorders in the kidney. During the past year, progress has been made in several directions. The theoretical study of qualitative simulation has been completed. QSIM, a new version of the qualitative simulation algorithm based on the theoretical work, has been implemented, runs very efficiently and is being distributed to interested researchers. The Evoker, an associative hypothesis-generator, has been redesigned and reimplemented to make it a more general-purpose tool for knowledge engineering. A preliminary version of the interaction between QSIM and the Evoker that constitutes the core of the RENAL program has been

designed and implemented. The size and scope of the knowledge bases for both QSIM and the Evoker are steadily being augmented. This project focuses on completing the knowledge bases, codifying a set of principles for acquiring and validating its contents.

The pattern of interactions are being completed between QSIM and the Evoker, concentrating particularly on the matcher between the different descriptive terms used by the two reasoners, to make RENAL into a complete diagnostic system.

Software Engineering

Software Studies

Yale University; Elliot M. Soloway; *Mapping Between Programmers' Conceptualizations And Programming Language Constructs (Computer Research and Information Science)*; (DCR-8302382 A02); $13,711; (Joint support with the Division of Information Science and Technology - Total Grant $19,066).

Expert programmers have and use rich mental conceptualizations that capture information beyond the syntax and semantics of programming languages, such as when to use a construct, how to integrate one construct with another. The hypothesis being explored is: "If a programming language construct mirrors a programmer's conceptualization of the task, then that construct can be used more effectively than one that does not." Thus, the mappings between programmers' mental models and programming language constructs, for 3 basic programming notions: variables, conditionality, and looping are being examined.

Three empirical techniques are being employed:

1. In construct comparison studies a construct that matches a programmer's mental model better than an existing language construct is being created; it is conjectured that a group using the new construct should out perform (on accuracy, speed) a group using the existing construct. Tasks being considered are program generation, modification, and debugging.
2. Bug patterns and frequencies, and
3. Video-taped interviews are being used as corroborating evidence.

Examining how programmers comprehend specific problem solving and programming notions, and how they realize these notions in programming languages should provide important information to the designers of programming languages, environments, and methodologies.

Purdue University; Vincent Shen, Samuel Conte, and Hubert Dunsmore; *Software Metrics For Large-Scale Program Development (Computer Development)*; (DCR-8411699); $104,996; 24 mos.

Independent studies of existing software development models show that all of them perform poorly when applied to data bases of projects other than those used in developing the models. A major problem associated with these models is that the model coefficients reflect the environment in which they were developed, but may not be transportable to other environments. In addition, when such models are used for effort prediction, they rely heavily on an estimate of size (usually in lines of code). Early size estimation is a problem of comparable difficulty to that of early effort estimation. Methods for addressing each of these problems are being investigated. In particular, this research is investigating a technique for accurate early size estimation based on an estimate of a data structure metric which should be available early in the life cycle. The research is also investigating and extending the recently-developed COPMO (COoperative Programming MOdel) which is capable of producing acceptably-accurate environment-independent effort estimates based on some reasonably-defined complexity classes. The validation of these models is based partially on controlled experiments and partially on data from industrial sources.

University of Massachusetts - Amherst; Edward M. Riseman and Arthur S. Gaylord; *A Group Research Facility for Artificial Intelligence, Distributed Computing, and Software Systems (Computer Research)*; (DCR-8318776 A02); $25,000; 12 mos; (Joint support with the Special Projects, and the Intelligent Systems Programs - Total Grant $120,000).

This experimental computer research project focuses research in six areas: computer vision, robotics, distributed computing, software development, intelligent interfaces, and natural language processing. While most of the actual research activities are supported under other grants and contracts, including some from NSF, this grant provides partial support for the experimental research facility used by all the researchers, supplementing institutional, departmental, and other funds. NSF support is used for administration and operation of the COINS research facility, with the funding going towards technical support staff and equipment maintenance.

North Carolina State University - Raleigh; William Stewart; *Approximate Solutions of Queueing Models — An Error Analysis (Computer Research)*; (DCR-8503067); $30,881; 12 mos; (Joint support with the Computer Systems Design Program - Total Grant $45,881).

This research investigates the error behavior of iterative approximation methods. Phase I characterizes iterative methods as systems of non-linear equations and also categorizes the problems to which the iterative methods have been applied to develop a comprehensive testbed of models. Phase II is a systematic study of

the errors generated by iterative methods. The study involves implementing the different methods and generating results for the models in the testbed over a range of parameter values. Efforts are being made to develop mathematical expressions for the errors.

University of Utah; Robert E. Barnhill; *Surfaces in Computer Aided Geometric Design (Computer Research)*; (DCR-8502858); $45,000; 24 mos; (Joint support with the Computer Systems Design Program - Total Grant $177,431).

This research deals with the creation of new surface forms and the determination of geometric properties of surfaces in an interactive computer graphics environment. The investigations of new surface forms include the topics of Bernstein-Bezier triangular interpolants, surfaces defined on surfaces, and sweep surfaces. The geometric properties considered are estimation of derivative data, visual (or geometric) continuity and reflection lines of surfaces. In creating surfaces, the geometric data given are "shape" information given either as "hard" data such as positions or as "soft" data such as a designer's conception of a pleasing surface. In typical problems one encounters a mixture of both types of data. The surfaces to be created must therefore satisfy a variety of criteria such as sufficient smoothness, reasonable shape, and sufficient accuracy or fidelity. These requirements cannot all be met by the standard tensor product surfaces often used in practice. This lack is due to the rigid rectilinear structure required by tensor product schemes. Consequently, the research takes the more ambitious approach of developing a robust class of surfaces that can also utilize arbitrarily located information. The research involves the mathematics of surfaces, algorithms for the rendering of surfaces, the computer graphics implementation of these algorithms, and the recording of the visual results by means of some photographic medium. This research area is widely applicable in "CAD/CAM" (Computer Aided Design/Computer Aided Manufacturing), an engineering discipline dedicated to the automation of design and manufacturing.

Virginia Polytechnic Institute and State University; H. R. Hartson and R. W. Ehrlich; *The Structure of Human-Computer Interaction (Information Science and Computer Research)*; (IST-8310414 A01); $30,964; 12 mos; (Joint support with the Division of Information Science and Technology - Total Grant $123,855).

The goal of this research is to develop a theory of the dialog which takes place between an information system and its user and to embody that theory in a Dialog Management System, which then serves as a test bed for the theory's validation. The theory is based on a model which is being developed to provide a framework for discussion of all aspects of a user-system dialog without regard to the specific context of that dialog or to the specific technologies by which it might be implemented.

Virginia Polytechnic Institute and State University; Dennis G. Kafura; *Validation and Application of Software Metrics to Design and Maintenance (Computer Research)*; (DCR-8418257); $89,966; 24 mos.

To make software an engineerable product, it is important that the designers, implementors, and maintainers of software systems be able to express the characteristics of the system in objective and quantitative terms. One technique for producing the needed metrics is based upon an automated analysis of the structure of a software system. Over the last seven years Professor Kafura has identified a collection of structure metrics and has reported encouraging experience in their use. This research continues past work by:

1. Performing additional quantitative validations of production systems from NASA/Goddard; and
2. Investigating the use of structure metrics in guiding and controlling the system design and maintenance process.

Program Testing and Fault Tolerance

University of California - Irvine; Nancy G. Leveson; *Analyzing Fault Tolerance and Safety Using Time Petri Nets (Computer Research)*; (DCR-8406532); $84,554; 24 mos.

The problems of safety have become important as computers are increasingly being used in applications where the consequences of failure are serious and may involve grave danger to human life and property. This project involves exploring the application of Time Petri net modeling and analysis techniques to safety-critical real-time system design. The goal is to develop analysis procedures to identify hazards and safety-critical single and multiple failure sequences, to determine safety requirements including timing requirements, and to analyze the design for fault tolerance and safety.

University of Massachusetts - Amherst; Lori A. Clarke, Debra J. Richardson, and Steven J. Zeil; *The Development of a Scientific Testing Method*; (DCR-8404217 A01); $119,238; 12 mos.

There are a number of research projects addressing the problem of software testing. For the most part, recent research has moved from developing tools that gather information about programs to developing tech-

niques that actually apply this information. There has been little work, however, on understanding the comparative strengths and weaknesses of existing techniques or their potential interactions. In this project several testing techniques are being evaluated to gain a better understanding of their strengths, weaknesses, and overlap. This evaluation should lead to insights on how to extend and improve some of these techniques and how to integrate a set of these techniques into a powerful testing method. The goal is to assure that the error detection capabilities of this integrated method are well-understood and clearly stated.

Cornell University; Kenneth Birman; *Research in Fault-Tolerant Distributed Computing (Computer Research)*; (DCR-8412582); $27,978; 12 mos.

This project investigates fault-tolerance and performance in distributed computer software, a topic of importance and interest. "Resilient objects", which reside at independent processing sites replicating data and synchronizing actions, are used to ensure that correction constraints are respected. Using "Resilient objects", the system rather than the programmer is concerned with managing replicated data and controlling interactions between sites. This research concentrates on theoretic aspects of maximizing concurrency in fault-tolerant distributed systems.

New York University; Elaine Weyuker; *The Use of Data Flow Information to Select and Evaluate Test Data and Evaluate Software Complexity (Computer Research)*; (DCR-8501614); $45,500; 12 mos.

This research focuses on the investigation of the use of data flow information in selecting software test data, evaluating software test data, and evaluating the complexity of software. These investigations involve constructing a data flow analysis tool, extending the theory of testing criteria, assessing the cost of using the proposed criteria, evaluating their effectiveness, defining a software complexity metric, and evaluating this metric.

Software Tools and Programming Environments

University of California - Berkeley; Lawrence A. Rowe; *A Program Development Environment for Interactive Database Applications (Computer Research)*; (DCR-8507256); $73,988; 12 mos.

This research focusses on program development environments for interactive database applications. The research includes the investigation of graphical views of an application for program browsing and specification, rapid development of office applications, and spreadsheet interfaces to databases. This work is being carried out in the context of a new programming environment based on filling in forms to specify an application.

University of California - Irvine; Peter A. Freeman; *Reusable Software Engineering (Computer Research)*; (DCR-8304439 A02); $119,322; 12 mos.

The goal of this project is the development of effective reusable software engineering (defined as activities that produce desired software engineering workproducts and information that can be reused later in producing a new workproduct). This work is a continuation and expansion of current work and falls into two categories of experimentally based work:

1. Software construction using components; and
2. Advanced models and tools.

The first category is a continuation of work on the Draco approach to the use of software components that utilizes program transformations and domain-specific languages. The second category addresses the key problem of representation of system development information and is exploring the reuse of graphical models of system design information.

University of California - Irvine; Richard Taylor; *Arcadia: An Environment for the Analysis, Testing, and Debugging of Concurrent Software*; (DCR-8502558); $47,497; 24 mos.

Assuring that real-time or concurrent software performs its intended function is a difficult task, arguably more so than for serial software. Moreover if such software is developed on one machine, but final execution occurs on another, additional difficulties arise. Recent research has developed some promising new techniques for aiding in the checkout of this type of code, including static concurrency analysis, program instrumentation with run-time monitoring, and multitasking debugging techniques. Not surprisingly, however, no single technique is adequate for all the demands. What is necessary is their cooperative, mutually supportive application. The primary goal of this project is the further development of these and other techniques, emphasizing integrated application. An important recognition guiding the research is that effective investigation requires that the tools embodying the techniques be housed in a programming environment. In so doing duplication of tasks can be avoided, a uniform user interface can be provided, and

efficiency can be promoted. This project's analysis, testing, and debugging tools are being developed in concert with a programming environment. This environment, a platform for research, draws heavily from work with the Toolpack and Arcturus environments, and emphasizes extensibility and flexibility.

University of Southern California; Robert M. Balzer and David Wile; *Transformation-Based Maintenance (Computer Research)*; (DCR-8304190 A02); $111,500; 12 mos; (Joint support with the Software Systems Science Program - Total Grant $223,000).

For several years this project has been creating the technology needed to support an alternative software development paradigm. This paradigm is based on capturing the entire program development process - including specification, design, implementation and maintenance - within the computer and suppporting it with automated tools. The investigators believe this paradigm will ultimately result in dramatic improvements in both software productivity and reliability because the human developer's role will be reduced to decision making and guidance, while the computer's will be expanded to manipulation, analysis and documentation. To this end, the primary research focus has been on recording, structuring, and formalizing the program development process, making it amenable to machine aid and manipulation.

The investigators expect that the leverage provided by this paradigm will derive more from its support of the maintenance process than from its support of the initial implementation, wherein the increased formalization of the development process requires more involvement by the human developer. On the other hand, maintenance involves small changes. Increased formalization of these changes is insignificant if the already formalized development can be adapted and reused. Therefore, the project focuses on supporting such reimplementations by strengthening the formal representation of developments, by supporting the evolution of specifications through the same mechanisms that support maintenance, and by developing a formal language of modification for both specifications and developments.

Rocky Mountain Institute for Software Engineering; William E. Riddle; *Software Engineering Environment Evolution Workshop, Aspen, Colorado, November, 1985*; (DCR-8513620); $6,900; 2 mos; (Joint support with the Software Systems Science Program - Total Grant $9,800).

This workshop focuses on the development of a plan for research and development activities which further develop current software engineering environments to meet projected needs in the mid-1990's. Issues of primary focus are: integration, extensibility, and distributed architectures. It is expected that the gradual maturation of a sequence of prototype environments will be a major theme of the plan.

The four-day, by-invitation-only workshop was devoted to expanding and refining a draft plan prepared by the organizers. The plan will be exposed to the community at large through reports, conference panels, and journal articles as appropriate. The workshop-developed plan will help focus the efforts of the software engineering environments research and development community over the next five years.

University of Colorado - Boulder; Leon Osterweil; *Evaluation and Continued Developments of Software Tools and Environments*; (DCR-8403341 A01); $162,000; 12 mos.

This research centers on the continued development and evaluation of tools and environments to support the production of software. The research builds upon work done in the past few years on the Toolpack project. This project has resulted in the creation of a set of tools and integration software capable of providing considerable support for software development, maintenance, and testing. The architecture of the tools and integration software has been carefully designed to accommodate additions and changes to the toolset and command language. Thus this software seems to provide an excellent testbed for experimentation with tools, as well as substantive experimentation with the flexibility and extensibility of the Toolpack architecture itself.

University of Massachusetts - Amherst; Lori A. Clarke and Jack C. Wileden; *A Systematic Treatment of Interface Control (Computer Research)*; (DCR-8408143); $150,626; 24 mos.

Although the importance of interface control in the production and maintenance of large, complex software systems has been repeatedly recognized, there has been no systematic treatment of interface control and its role in the software development process. Existing languages provide incomplete and imprecise mechanisms for indicating which entities in a software system can access other entities and how they are allowed to manipulate those entities. Moreover, these weaknesses undermine the effectiveness of any analysis that can be done to assure consistency among interfaces and proper use of shared entities. This research, continuing preliminary investigations of a language framework to support interface control, focuses on:

1. Further development and use a formal model to evaluate their language framework;
2. Incorporation of the framework into a representative set of specifications, design, and programming languages; and
3. The design and construction of a prototype set of analysis tools to support the use of the framework.

Cornell University; Robert L. Constable; *Experiments with a Program Refinement System*; (DCR-8303327 A02); $63,494; 12 mos; (Joint support with the Software Systems Science Program and the Intelligent Systems Program - Total Grant $189,494).

This research involves the judicious combination of the strong points of program development methodology and program verification in a logic for correct program development. This project is exploring a logic for formal top-down development called refinement logic. A system called Program Refinement Logic (PRL) has been designed and implemented. This system is being improved and expanded and considerable experimentation is being conducted.

Rice University; Kenneth Kennedy; *Interprocedural Information in a Programming Environment (Computer Research)*; (DCR-8503608); $90,363; 12 mos.

The R^n Fortran Programming Environment provides an integrated collection of tools to aid programming teams in entering and testing Fortran programs, along with sophisticated facilities for managing all the code in the various program components. An important part of the R^n environment is the project data base, in which is recorded all the information that the environment knows about the programs and modules that comprise the project.

This investigation is examining ways that interprocedural information, automatically recorded in the data base by the various tools, can be used to facilitate the process of program design, and construction. This work concentrates on the roles of the editor, debugger, program composition processor, optimizing compiler, and execution manager in this information exchange.

University of Virginia; Terrence Pratt; *A Programming Environment Based on Formal Semantic Models (Computer Research)*; (DCR-8502881); $46,662; 12 mos.

This research investigates the construction of a programming environment based upon a unifying formal model, H-graph semantics. All analysis, verification, and transformation tools provided by the environment work from the H-graph models of programs. The formal models include the entire system (application software, language support software, operating systems, and hardware).

University of Washington; David S. Notkin; *Customization and Uniformity in Flexible Systems (Computer Research)*; (DCR-8505859); $44,988; 12 mos.

Customization occurs when users alter the appearance or performance of a system by selecting from among a range of choices predetermined by the designer. Because the design and implementation of customizations is left to individual system designers, most environments lack consistency among their subsystems. This research is concerned with developing models and implementations of powerful customization mechanisms that support uniformity throughout a computing environment.

The model of customization is based on a customization library that is shared by all system designers in the environment. By centralizing the customizations available in the environment, uniform system development is facilitated. The implementation of this model uses both the centralized library and also multiple customization profiles, which define the customization values selected by each user. Both the library and the profiles are declarative, non-procedural notions that are easy to manipulate and modify (both for implementors and for users, including non-technical users). These mechanisms additionally lend themselves to straightforward, efficient implementation. The models, mechanisms, and implementation lay the framework for development of uniform, flexible environments that permit more complex modifications of their underlying systems.

Compilers and Language Implementations

University of Arizona; Ralph E. Griswold; *The Implementation of Generators and Goal-Directed Evaluation (Computer Research)*; (DCR-8502015); $65,596; 12 mos.

Generators and goal-directed evaluation, combined with novel control structures, provide an expression evaluation mechanism in the Icon programming language that allows particularly concise and natural formulations for many complex programming language problems. This expression evaluation mechanism is potentially applicable to a variety of other programming languages, such as C, Pascal, and Smalltalk.

Efficient implementation of this expression evaluation mechanism is difficult and there is presently no coherent model that can be used as a basis for implementations in new programming language contexts.

This proposal requests support to evaluate the present implementation of expression evaluation in Icon and to develop a more efficient and portable approach to new implementations. This research will interact with research on new programming language features, leading to a better understanding of the relationships between novel aspects of expression evaluation and their implementation.

University of Arizona; Ralph E. Griswold; *High-Level Programming Language Facilities for Data Structure Processing*; (DCR-8401831 A01); $80,364; 12 mos.

The focus of this investigation is the study of high-level programming language facilities for processing nonnumerical data, such as strings and lists. Previous work has concentrated on two efforts:

1. The design of new facilities for characterizing and processing nonnumerical data, and
2. The implementation of these facilities so that they can be evaluated and made available to the computing community.

There is a close interaction between these two efforts: the value of new facilities often can be determined only by their application to real problems. Conversely, the use of new facilities frequently suggests new possibilities and neglected areas of application.

This research builds on past results and focuses on two areas of investigation:

1. The development of a laboratory for the study and design of high-level pattern-matching facilities for strings, structures, and data bases, and
2. The design of programming language facilities for manipulating sequences as data objects.

University of Arizona; Eugene W. Myers; *The Design and Analysis of Efficient Applicative Data Types and Their Applications (Computer Research)*; (DCR-8511455); $47,943; 12 mos.

This research addresses issues associated with the efficient realization of applicative paradigms on von Neumann architectures. Ways to improve the underlying designs of the algorithms and data structures that realize applicative data types are being investigated. Specific applications of the technology in programs such as debuggers are being developed.

Colorado State University; Robert Mueller; *Symbolic Covers in Global Retargetable Generation and Compaction (Computer Research)*; (DCR -8503941); $87,489; 24 mos.

This research continues Mueller's work on the development of a theory and methodology for automatically retargeting microcode. Mueller has been successful in the areas of retargetable code selection and local compaction and is now seeking a general framework for global retargetable microcode generation and compaction. A phase-coupled model based on the symbolic cover of Reif is being explored.

Brandeis University; Jacques Cohen; *Compiler Prototyping Using Logic Programming (Computer Research)*; (DCR-8500881); $42,982; 12 mos; (Joint support with the Intelligent Systems Program - Total Grant $66,982).

This research is applying techniques developed using Prolog to current topics in compiler development, including parallel compilation and code generation. Equally important is the investigation of the development of efficient compilers and interpreters for extensions of Prolog itself. The goal of the research is to produce Prolog programs for compiler applications that will execute as efficiently as typical C programs.

University of Utah; Robert R. Kessler; *Implementation Strategies For Portable LISP Based Systems (Computer Research)*; (DCR-8204247 A04); $132,804; 12 mos.

This project centers on the portable standard LISP system (PSL). Past work has led to the development of the PSL on at least six different systems ranging from a Motorola 68000 to a Cray-XMP. Recent work has focussed on design and initial implementation of a better portability strategy, development of a peephole optimizer based on an architectural description language and the development of a system for analyzing PSL source code. The primary goal of this research is the development of a new LISP compiler which will provide better portability, better modularity, support for data typing, graph coloring based register allocation and improved code generation.

University of Washington; Robert Henry; *Machine Dependencies in Retargetable Compilers (Computer Research)*; (DCR-8501818); $49,659; 12 mos.

This research pursues three interrelated lines of work to study, improve, and quantify the effects that machine dependencies have upon retargetable compilers. New tree matching algorithms are being applied to code generation; phase ordering problems involving the interaction of machine independent and machine specific optimizations are being investigated; and the effect of compiler design on architectures designed in conjunction with compilers is being measured.

Software Systems Science

Programming Languages: Design, Semantics, Implementation

University of Southern California: Edward K. Blum; *Algebraic Approaches to Programming Concepts and Their Semantics (Computer Research)*; (DCR-8406920 A01); $62,731; 12 mos.

The main objective of this project is to analyze and develop the relations connecting operational semantics to algebraic semantics of programming concepts and programming languages. This entails a study of the applicability of algebraic and more general axiomatic methods to the semantics of programming languages and systems. One area of emphasis is the implementation of data types by an algebraic method recently initiated by the investigators. This method is being developed further and its practicality is being tested on the I.E.E.E. floating-point type and the proposed Fortran array type, both of which provide real test cases for studying the relation of algebraic to operational semantics. A second area of emphasis is the proper mathematical framework for recursive derived operations which arise in implementations of data types and general computations over algebras. A third area of emphasis is the algebraic and operational semantics of concurrent systems.

University of Illinois - Urbana: David J. Kuck and Duncan H. Lawrie; *Software Studies for Multiprocessors and Memory Hierarchies*; (DCR-8406916); $148,458; 12 mos.

The central theme of this project is the exploitation of parallelism in multiprogrammed multiprocessors and the management of memory hierarchies of multiprocessors and monoprocessors with large main memories. The existing Parafrase system for restructuring Fortran programs has been quite successful in speeding up computations on array-type or parallel processors and in reducing space-time products in memory hierarchies. The present goal is to find approaches for dealing with the remaining small percentages of serial running time that Parafrase cannot transform to parallel code. A multiprogrammed, multiprocessed system simulator is being implemented to enable measurements and experiments to be performed with different scheduling and resource allocation strategies.

Indiana University - Bloomington: Daniel P. Friedman and Christopher T. Haynes; *Liberating and Constraining Control (Computer Research)*; (DCR-8501277); $94,237; 12 mos.

Many programming languages provide control objects, such as Algol 60 labels, but such objects are rarely first-class. Continuations, though originally introduced as meta-linguistic abstractions, may be provided as first-class objects for direct use by programmers — as in the language Scheme. Given access to continuations, it is possible to extend a language to include any desired sequential control abstraction. Thus continuations are a very powerful tool for language extensibility. However, in specific applications it is often desirable to place constraints on the use of continuations. Such constraints may allow more efficient implementation, aid in reasoning about programs, and help enforce security conditions.

This project is directed at better understanding the effects of providing access to first class continuations and studying the types of constraints that may be required and how they may be enforced.

University of Kansas: Adrian Tang and Tsutomu Kamimura; *Finitely Continuous Posets as Domains (Computer Research)*; (DCR-8415919); $104,241; 18 mos.

This project is investigating the notion of finitely continuous posets as the unifying notion of domain for denotational semantics of modern programming languages. Finitely continuous posets include all the existing notions of domains and are closed under various domain constructions. It is conjectured that the category of finitely continuous posets is indeed the largest cartesian closed category of continuous domains. Finitely continuous posets are being studied in three directions. First, an attempt is being made to characterize finitely continuous posets using projective sequences of some finite structures. Such a characterization would provide an intuitive formulation of the notion of an effectively given domain. Second, universal domains for finitely continuous posets are being investigated. In particular, it is important to know whether the universal domain, called FP, is indeed a natural model for studying bounded nondeterminism. An effort is being made to obtain an intrinsic characterization of the computable retracts of FP. Finally, the possibility of some nonalgebraic behavior is being studied and whether such construct can be handled in the framework of finitely continuous posets is being examined.

Massachusetts Institute of Technology; John V. Guttag; *Industry/University Cooperative Project: Extensions to the Rewrite Rule Laboratory and Development of Related Theory (Computer Research)*; (DCR-8411639); $25,958; 12 mos; (Joint support with the Theoretical Computer Science Program and the Division of Industrial Science and Technological Innovation - Total Grant $103,831).

This project is focused on three related topics:

1. Enhancement, maintenance and distribution of rewrite rule laboratories - REVE and RRL,
2. Theoretical and implementation work on the method of proof by consistency, and
3. Theoretical and implementation work on E-unification and generalizations for the Knuth-Bendix procedure.

University of North Carolina - Chapel Hill; Richard T. Snodgrass; *The Formal Specification and Implementation of a Temporal Query Language*; (DCR-8402339 A01); $76,254; 12 mos.

The goal of this project is two-fold: to develop a formal semantics for a query language for relational databases containing historical information, and to implement a historical database management system utilizing this language. The tuple calculus semantics, already defined for a subset of the language is being extended to handle aggregate functions and incomplete information. The implementation is based on a temporal relational algebra that supports the incremental updating of historical relations. This algebra is being formalized and proved to preserve the semantics defined by the tuple calculus. Optimization strategies for historical databases are also being investigated.

New Mexico State University; Mark B. Wells; *Development of a Modcap System (Computer Research)*; (DCR-8303931 A01); $116,584; 12 mos.

Modcap is a programming language based on the fundamental idea that (1) functions be first-class objects (in a block-structured language) and (2) primitive and user-defined data types naturally form a lattice of sets. Full exploitation of these concepts has produced a simple and general language out of which a user can build special-purpose applications-oriented programming systems whose use entails little loss of efficiency. This project involves:

1. The development of workable Modcap implementations for academic and corporate environments; and
2. Studies of associated programming methodology.

City University of New York - City College; Valentin Turchin and Robert Nirenberg; *A Supercompiler System Based on the Language REFAL (Computer Research)*; (DCR-8412986); $175,671; 24 mos.

This project will continue work on the programming system based on a supercompiler for REFAL. The purpose of this system is to facilitate the creation of new application-oriented programming languages at low expense, and to perform in the computer a great deal of work on optimization of algorithms, and even on the algorithmization itself, which is now performed by man. The user of the system will have to formulate only the definition of his problem, its mathematical model, without bothering about the details of algorithmic efficiency and data structure in the real computer. The core of the system is a program called a "supercompiler", which transforms a multi-level semantic definition of a problem in REFAL into an efficient computer program. It can also be used for program optimization, program investigation, and problem solving of different kinds including theorem proving. At present, a REFAL system which includes an experimental model of the supercompiler and a number of processors for intermediate languages exists.

Cornell University; Alexandru Nicolau; *Percolation Scheduling: A Hierarchical Parallel Compilation Technique (Computer Research)*; (DCR-8502884); $84,879; 24 mos.

Percolation Scheduling (PS) is a new technique for compiling programs into parallel code. It attempts to overcome problems that limit the effectiveness and applicability of currently available techniques.

PS globally rearranges code past basic block boundaries. Its core is a small set of primitive program transformations defined in terms of adjacent nodes in a program graph. These transformations are easy to understand and implement and are independent of any heuristics. They constitute the lowest level in a hierarchy of transformations and guidance rules. Higher levels of this hierarchy direct and enhance the applicability of the core transformations and enable the exploitation of parallelism, from the fine grained to the coarse. Aided by the other levels, the core transformations operate uniformly and can be applied to partially parallelized code. This allows modification of code produced by PS and improvement of code produced by less-effective compiliers. In addition, these transformations lend themselves to highly parallel execution, which may significantly reduce compilation time when run on a parallel machine.

General Electric Research and Development Center; David R. Musser; *Industry University Cooperative Research: Extensions to the Rewrite Rule Laboratory and Development of Related Theory (Computer Research)*; (DCR-8408461); $12,571; 12 mos; (Joint support with the Theoretical Computer Science

Program and the Division of Industrial Science and Technological Innovation - Total Grant $50,284).

This project is focused on three related topics:

1. Enhancement, maintenance and distribution of rewrite rule laboratories - REVE and RRL,
2. Theoretical and implementation work on the method of proof by consistency, and
3. Theoretical and implementation work on E-unification and generalizations for the Knuth-Bendix procedure.

The MIT/GE collaboration in devising a rewrite rule laboratory has made very good progress, to the point that they are nearly ready to distribute a coherent system for use by a small group of other rewrite rule researchers. This distribution, with the accompanying feedback and maintenance, is a milestone in this particular research area. Availability of a common laboratory system facilitates comparison of algorithms and logical procedures and provides a base for further work on proof by consistency.

State University of New York - Buffalo; Jeffery I. Zucker; *Denotational Semantics of Programming Languages (Computer Research)*; (DCR-8504296); $75,852; 12 mos.

This project continues denotational semantics and program verification in two directions: (1) To develop, and compare, various semantic frameworks, both operational and denotational, for programming languages involving concurrency, with merge, synchronization and communication. This includes (in the denotational case) linear time and branching time semantics, where, in the linear time case, both metric and order-theoretic structures are considered, and, in the branching time case, the theory of processes, developed by J.W. de Bakker, has a central role. (2) To write a research monograph (with J.V. Tucker) which develops a theory of program correctness over abstract data structures, with error-state semantics. This involves developing a theory of computability over abstract data structures.

Syracuse University; John C. Reynolds and Nancy J. McCracken; *The Design, Definition, and Implementation of Programming Languages*; (DCR-8017577 A05); $88,664; 12 mos.

Research is being carried out on the design of programming languages and the mathematical definition of their semantics, with emphasis on the design of Algol-based languages and the use of category theory as a definitional tool. Specific topics include: the application of category-sorted algebras to treat implicit conversions and generic operators; the use of a non-functional Cartesian closed category to provide an abstract model of the store; extension and modeling of a language in which interference can be determined syntactically, and development of a variant of specification logic for proving programs in such a language; and models of languages with type binding.

Carnegie-Mellon University; Edmund M. Clarke; *Design and Verification of Concurrent Systems (Computer Research)*; (DCR-8216706 A02); $30,452; 12 mos.

The rapid development of computer networks and multiprocessor systems has increased the need for methods of constructing reliable concurrent systems. This project is directed at improved understanding in three particularly important areas of concurrent system development:

1. Specification and verification of concurrent systems. Proof rules for liveness properties of concurrent programs are difficult to formulate because such properties lack continuous fix-point characterizations. It may be possible to construct concurrent programs in which the synchronization skeleton is automatically synthesized from a high-level temporal logic specification.
2. Construction of distributed algorithms. Protocols, such as the distributed termination protocol, that can be superimposed over a distributed program to ensure some particular behavior are being studied. Also, it may be possible to transform the specification of a sequential program into that of an equivalent distributed program.
3. Fault-tolerant computing. Multiple resources potentially enable a system to continue functioning despite hardware errors. Software to detect and to reconfigure the system into a reliable subsystem is required. A current line of research on such reconfiguration strategies is being extended.

University of Utah; Gary Lindstrom; *Generalizations of Functional and Logic Programming (Computer Research)*; (DCR-8506000); $81,659; 12 mos.

Functional and logic programming are two forms of what may be called assertional programming languages. This new generation of languages is characterized by discipline over side-effects, highly developed notions of program modularity, clean mathematical bases, and facilitation of concurrent execution. While functional and logic programming languages are similiar in many respects, they are also substantially different in conceptual and practical power. This project seeks to elucidate these differences, and amplify their combined power through the development of a more comprehensive framework for assertional programming. The specific goals of this research are to:

1. Amalgamate functional and logic programming into a new class of languages generalizing both;
2. Design an implementation strategy for such new languages, suitable for multiprocessor architectures; and
3. Demonstrate the added power of this class of

languages and implementations in two applications, namely a distributed rule-based programming system and a system for dynamic constraint maintenance (i.e., homostatic programming).

Operating Systems and Concurrency

New York University; Dennis Shasha; *Research on Semantically-Based Concurrency Control for Data Structures (Computer Research)*; (DCR-8501611); $47,817; 12 mos.

The data structure that implements an abstract data type (ADT) often contains more detail than required by the ADT. For example, a B-tree that implements a dictionary ADT not only contains the keys of the dictionary, but is also a balanced tree. Concurrent algorithms can exploit this difference in semantic detail to achieve more concurrency than syntactic concurrency control methods would suggest. This approach is being used to design and verify concurrent algorithms on data structures supporting various set and multi-set ADT's such as those used in databases, ADT's used by graph manipulation languages such as Lisp, and ADT's such as directories used in distributing computing environments. Recovery, performance analysis, and language issues related to this approach are also being examined.

State University of New York - Stony Brook; Arthur J. Bernstein; *Asynchronous Algorithms and Concurrency Control (Computer Research)*; (DCR-8502161); $68,168; 12 mos.

The project is primarily concerned with distributed algorithms, particularly as they relate to concurrency control. The starting point is work currently in progress to develop a concurrency control algorithm for query dominant systems in which read only transactions require no synchronization. A multi-version optimistic scheme has been developed for this purpose and the proposed research is connected with this effort:

1. The current algorithm is being extended to the distributed case, its correctness proved, its performance evaluated in the presence of failure and in comparison with alternative strategies.
2. Techniques are being studied for accessing a shared data base in which readers do not have to synchronize with concurrent writers (writers must, of course synchronize with each other). Some algorithms of this type already exist and are useful in eliminating synchronization overhead for read only transactions.
3. Applications are being examined which do not require serializability as a concurrency control constraint and techniques are being studied for detecting such situations and developing appropriate weaker constraints.

Additionally, work continues in multicast communication.

State University of New York - Stony Brook; Scott A. Smolka; *Livelock, Lockout, and Liveness in Networks of Communicating Finite-State Processes (Computer Research)*; (DCR-8505873); $39,015; 12 mos.

The problems of livelock, lockout, and liveness (L^3 for short) are being analyzed in the context of networks of finite-state processes. The effort concentrates on the issues of decidability, complexity, and design of efficient decision procedures. The investigation is based on a new algebraic technique for efficiently solving the problems of potential blocking and termination. Extending this technique to the L^3 problems, which are inherently more complex than the blocking and termination problems, warrants a heuristic approach.

A companion goal of the research is to produce an automatic protocol validation system equipped with a graphical user interface. This system provides a flexible environment in which to embed decision procedures and also allows, in a "painless" way, experimentation with different heuristics for the L^3 problems. Finally, the system is being used extensively to validate concurrent programs and network protocols.

Data Bases

University of California - Berkeley; Alan J. Smith; *Automatic File and I/O System Design and Optimization*; (DCR 8202591 A02); $78,670; 12 mos.

This research includes four related topics:

1. Development and evaluation of algorithms for the migration of files between secondary and tertiary storage (medium and long term migration),
2. Development and evaluation of designs for cache disks and for the algorithms to manage them,
3. Development and evaluation of algorithms for the management of distributed file systems, and
4. Analysis of user behavior and workloads, primarily as they relate to file migration in a memory hierarchy and across a distributed file system.

This research is developing a body of empirically validated algorithms and designs, based on theoretical and experimental studies, which significantly improve the operation of computer systems.

The methodology of this research is characterized as experimental computer science. Traces of real file system operation (most of which already exist and have been reduced to useable form) are being analyzed, characterized, and used as the basis for algorithm development and device design. The results of this research are being implemented on one or more DEC VAX 11/780's. Evaluation of the results of this research is being done both by implementation and by trace driven simulation. The results of this research should be applicable to a very wide range of computer installations from large minicomputers to massive computer centers. The wide variety of traces available permit tests of the range of applicability of the results.

University of Kentucky; Wiktor Marek; *Investigation on the Distributed Databases and Incomplete Information in Databases (Computer Research)*; (DCR-8411600); $82,752; 24 mos.

This project is investigating aspects of the theory of databases including the theory of distributed databases and the consistency problem for "future states" of databases. The relationship of the latter to the nonmonotonic logics of McDermott and Doyle is being investigated.

University of Massachusetts - Amherst; David W. Stemple; *Theoretical Support for Database System Development (Computer Research)*; (DCR-8503613); $208,757; 24 mos.

Database systems, like all models, must be constrained to represent just those states and transitions which are possible in the world they model. Database integrity constraints, transition constraints and transaction definitions specify the conformity of a database system to the real world. However, the implemented system must obey these specifications if the model is to be valid. The enforcement of database system constraints is a difficult problem to solve efficiently; this follows from the large amounts of data involved and the complexity of determining minimum required checks. One approach to achieving efficiency is for transactions to be structured in a manner which precludes the possibility of their disobeying any constraints. Such transactions are termed safe. Verifying the safety of a transaction allows it to be run without any constraint checks beyond those in the transaction itself. There is also promise that attempts to verify transaction safety can lead to information needed to transform unsafe transactions into efficient safe transactions.

This project is investigating techniques for automatically verifying the safety of transactions and transforming unsafe transactions into safe transactions. An effective approach to this problem involves the integration of formal analysis into a database system development method.

Cornell University; Dina Bitton; *A Methodology for Experimental Evaluation of Database Management Systems (Computer Research)*; (DCR-8410889); $108,442; 24 mos.

The goal of this project is to develop a performance evaluation methodolgy for database systems based on benchmarking. Besides identifying the strengths and weaknesses of a particular database system, benchmarks generated according to this methodolgy provide a much needed comparative tool, to evaluate the relative performance of database machines and software database management systems.

The approach radically differs from previous research on performance evaluation of database systems. Rather than performing application or machine oriented measurements, a general model of performance is being investigated, and a functional set of benchmarks that can systemically measure all the variables in this model is being developed.

New York University; Paul Spirakis; *Innovative Techniques for Concurrency Control in Database Systems (Computer Research)*; (DCR-8503497); $46,657; 12 mos.

This project concentrates on:

1. A semantic approach to correctness of concurrent execution of transactions;
2. The introduction of probabilistic techniques to the design of efficient concurrency control schedulers and protocols; and
3. The use of a highly parallel supercomputer architecture and software for correct, fast, concurrent execution of many transactions.

In (1) semantically correct schedules of transaction executions have been defined and research is being conducted on the discovery of classes of schedules (with suitable semantics of transaction actions and integrity constraints, e.g., linear semantics, statistical queries etc.) which are nonserializable but which are correct and have an easy membership test. In (2) the use of probabilistic bidding, random sampling and random waits, as means for designing efficient distributed concurrency control protocols is being studied. For example, locking methods can be treated as special cases of resource allocation problems, which can be solved efficiently by probabilistic techniques. In (3) the design of the algorithms and primitives to support a multi-user transaction processing environment, on top of a highly parallel multiprocessor system with shared memory, like the NYU Ultracomputer is being investigated.

State University of New York - Stony Brook; David S. Warren and Edward Sciore; *An Integrated Prolog-Database System (Computer Research)*; (DCR-8407688 A01); $94,553; 12 mos.

There is a close and fundamental relationship between the logic programming language Prolog and relational database query languages. Previous Prolog development has concentrated on efficient execution of relatively small recursive programs running entirely in memory. Relational database systems, on the other hand, have developed efficient ways to handle very large amounts of data stored on disk; they also handle multiple concurrent users, allowing update of shared data. This project involves research into how a single integrated system can efficiently perform all these functions.

The research includes the development of a prototype Prolog system that efficiently executes both recursive programs and retrievals from large databases stored on disk. It also includes theoretical work on how to include concurrency and database update cleanly and logically in the Prolog framework. The prototype system serves as a testbed for this theoretical work.

Case Western Reserve University; Z. Meral Ozsoyoglu and Gultekin Ozsoyoglu; *A Statistical Database Management System: Design, Implementation and Experiments (Computer Research)*; (DCR-8306616 A03); $68,430; 12 mos.

A Statistical Database Management System (SDBMS) provides statistical data analysis capabilities as well as capabilities to store, model and manipulate data suitable for statistical operations. The statistical database (SDB) security problem is to limit the use of the SDB in order to prevent users from inferring protected information in the database. An SDBMS should also have SDB security enforcement capabilities.

A preliminary design of a SDBMS, called a System for Statistical Databases (SSDB) has been recently completed. It has the following major features:

1. A semantic data model with tools to model data for statistical operations,
2. A screen oriented query language to manipulate summary and raw data, and
3. A general purpose SDBMS useful in a variety of applications with the capability to enforce SDB security mechanisms.

This research project is studying SDBMS's and consists of two parts. In the first part, the research team is studying SDB problems within the framework of the design and implementation issues of SSDB. The second part of the project consists of experimenting with a test-bed implementation of SSDB in order to evaluate efficiency, effectiveness and usefulness, and carrying out a comparative performance evaluation of the proposed SDB security mechanisms in real-world applications.

University of Texas - Austin; Abraham Silberschatz and Henry F. Korth; *The Theory of Non-First-Normal Form Relations (Computer Research)*; (DCR-8507224); $83,558; 12 mos.

Until recently, almost all research in relational database theory assumed relations to be in first normal form, that is all domains consist of atomic elements. This assumption served well in database systems designed for applications derived from data-processing problems. However, recent efforts to apply the relational data model to non-data-processing applications such as user interfaces, operating systems, CAD databases, and decision support have beeen impeded by the first-normal-form constraint. The purpose of this project is to extend the relational model and obtain formal results regarding the structure of non-first-normal-form relational languages and databases. Some initial promising results in this area have already been obtained. These results, form the foundation of the proposed work. The directions of the work include language design, query processing, dependency theory, normal forms,, and null values in the context of non-first-normal-form relations.

Programming Methodology

Rocky Mountain Institute for Software Engineering; William E. Riddle; *Software Engineering Environment Evolution Workshop, Aspen, Colorado, November, 1985*; (DCR-8513620); $2,900; 2 mos; (Joint support with the Software Engineering Program - Total Grant $9,800).

This workshop focuses on the development of a plan for research and development activities which further develop current software engineering environments to meet projected needs in the mid-1990's. Issues of primary focus are: integration, extensibility, and distributed architectures. It is expected that the gradual maturation of a sequence of prototype environments will be a major theme of the plan.

The four-day, by-invitation-only workshop was devoted to expanding and refining a draft plan prepared by the organizers. The plan will be exposed to the community at large through reports, conference panels, and journal articles as appropriate.

The workshop-developed plan will help focus the efforts of the software engineering environments research and development community over the next five years.

California State University - Hayward; Louise E. Moser; *RUI: Automated Deduction: Decidable Subtheories of Graph Theory (Computer Research)*; (DCR-8408544); $56,209; 12 mos.

This project is directed at the investigation and implemention of decision procedures for subtheories of graph theory. Such subtheories have application to the verification of the design and implementation of large computer programs and systems. The design of a program or system can be, and often is, modeled and visually displayed as a graph. This graph can depict the rigorous functional behavior of the program. More often the graphical representation presents a partial, incomplete view of the system, a view that is useful to its designers. Properties and attributes of the graph can be used to verify that the program or system will behave as intended, or that one graphical representation is a correct refinement of another, within the view of the system chosen by its designer. Available automated deduction systems typically operate in domains, such as predicate calculus or recursive function theory, that are useful for investigating the functional behavior of programs. Mechanical deductive support for graphical properties is not currently available. The research is identifying decidable subtheories of graph theory, investigating their utility in typical applications, and constructing prototype decision procedures.

SRI International; Richard J. Waldinger; *The Automatic Synthesis of Computer Programs*; (DCR-8512356); $92,851; 12 mos.

Program synthesis research attempts to use the power of the computer to aid in the programming process itself. The ultimate goal of this research is the development of an automatic program-synthesis system. The user of such a system would provide a specification, i.e., a description of the desired program's purpose; the system would then design a program, i.e., a detailed sequence of instructions, for achieving that purpose. The program would be guaranteed to satisfy the specification and so would require no further verification.

More immediate goals of the project include the development of conceptual and mechanical aids for the human programmer and methods for the verification, modification, and optimization of conventionally produced programs.

The particular emphasis of this project is on techniques that are based on automated deduction, or theorem proving. To construct a program is regarded as a task of proving a mathematical theorem. Once the theorem has been proved, the desired program is extracted from the proof. Special theorem-proving techniques are being devised to support the program synthesis application.

Stanford University; Donald E. Knuth; *Analysis of Algorithms (Computer Research)*; (DCR-8300984 A02); $201,602; 12 mos.

This research is oriented towards an improved quantitative understanding of the fundamental algorithms of computer science, and it comprises five major activities:

1. Development of new, efficient methods for use on computers;
2. Analysis of the performance of important computer techniques;
3. Analysis of the computational complexity of problems.
4. Studies of programming languages and oftware methodology; and
5. Combinatorial theory to support the above activities.

Stanford University; Zohar Manna; *Interactive Program Synthesis (Computer Research)*; (DCR-8214523 A02); $90,944; 12 mos.

The goal of research on program synthesis is to allow the problem-solving ability of the computer itself to assist in the programming task. A program synthesis system accepts from its user a specification which gives no hint of a method by which that purpose can be achieved; the design of the program is the responsibility of the system.

The special thrust of this project is the design and implementation of an experimental interactive system, which would carry out the routine details of the program's derivation, but be guided in its major decisions by a human collaborator. The techniques being used are based on the "deductive approach" of Manna and Waldinger, in which the problem of constructing a program is regarded as a task of proving a mathematical theorem.

University of Southern California; Robert M. Balzer and David Wile; *Transformational-Based Maintenance (Computer Research)*; (DCR-8304190 A02); $111,500; 12 mos; (Joint support with the Software Engineering Program - Total Grant $223,000).

For several years this project has been creating the technology needed to support an alternative software development paradigm. This paradigm is based on capturing the entire program development process - including specification, design, implementation and maintenance - within the computer and suppporting it with automated tools. The investigators believe this paradigm will ultimately result in dramatic improvements in both software productivity and reliability because the human developer's role will be reduced to decision making and guidance, while the computer's will be expanded to manipulation, analysis and docu-

mentation. To this end, the primary research focus has been on recording, structuring, and formalizing the program development process, making it amenable to machine aid and manipulation.

The investigators expect that the leverage provided by this paradigm will derive more from its support of the maintenance process than from its support of the initial implementation, wherein the increased formalization of the development process requires more involvement by the human developer. On the other hand, maintenance involves small changes. Increased formalization of these changes is insignificant if the already formalized development can be adapted and reused. Therefore, the project focuses on supporting such reimplementations by strengthening the formal representation of developments, by supporting the evolution of specifications through the same mechanisms that support maintenance, and by developing a formal language of modification for both specifications and developments.

Indiana University - Bloomington; David S. Wise and Steven D. Johnson; *Methods and Architectures for Applicative Programming (Computer Research)*; (DCR-8405241 A01); $82,818; 12 mos.

This project is advancing the study of applicative languages as a basis for general purpose programming. The research has two main aspects: the formulation of constructs and methods required for robust programming endeavors, and the construction of hardware that supports the tenets of programming being espoused. In previous investigations a small number of language constructs have been isolated that serve to specify a surprisingly broad class of applications, including those that commonly arise in systems programming. "Data recursion" and indeterminate constructor have been added to the applicative vocabulary. With data recursions cyclic behavior can be represented by infinite data structures. The indeterminate constructor is used to address concurrency and real-time behavior. These constructs are related. Each enforces the programmer's view of activity as an attribute of data and thus its treatment in spatial terms.

This research focuses on the individual programmer and yields methods that improve productivity. The crucial question is whether these methods work "in the large", for example, in applications involving many participants. To explore such issues an existing, purely applicative programming language is being enhanced into a programming system.

In its present form the language implements "suspending construction", an operational model of computation that was proposed in the mid 1970s. This model can serve as a basis for a multiprocessing host that supports efficient execution of applicative programs. A sequence of prototype architectures based on model is being explored.

University of Southwestern Louisiana; Subrata Dasgupta; *Formal Design and Verification of Firmware*; (DCR-8408750); $102,728; 24 mos.

One of the explicit goals of firmware engineering is to provide a basis for the formal and rigorous verification of microprograms. The practical importance of this issue is clear: since firmware constitutes one of the key lower levels in the multilevel structure of computer systems, any error in the microcode could have serious and costly repercussions on the reliability of the higher levels of the system. The primary problem in constructing a system for firmware verification lies in that microprograms are inherently machine-specific. Thus, the verification methods used must take into account such hardware issues as side-effects, low-level parallelism and timing. This project is investigating the formal design and verification of firmware based on:

1. A general model of micro-architectures and an associate architecture description language;
2. An axiomatic proof theory based on Hoare and temporal logics; and
3. A hierarchical design strategy for the formal design of firmware.

Princeton University; Robert Sedgewick; *Analysis and Implementation of Algorithms (Computer Research)*; (DCR-8514922); $49,173; 4 mos; (Joint support with the Theoretical Computer Science Program - Total Grant $49,173).

This research concerns the detailed study of the properties of fundamental algorithms. The research centers on drawing specific conclusions on the best way to solve important practical problems, backed up by careful implementations of basic algorithms, full mathematical analysis, and empirical studies to verify the analysis. At best, this type of careful scrutiny often leads to the design of new and more efficient algorithms or to the identification of general techniques for use in studying algorithms. The research contributes to the understanding of important methods of problem solution.

Typically, research activity in this field spans a wide range of areas, from topics such as number theory and complexity analysis to topics such as programming languages and computer architecture as they relate to a narrower range of algorithms and problem domains, which are studied in detail (always with an eye toward those techniques which are generally applicable). This type of detailed analysis is reserved for the most important problems. Specifically, the research is developing the mathematical analysis necessary to obtain detailed computational complexity results. This is done in the context of important sorting, searching, and computational geometry algorithms. The research also makes heavy use of tools that allow the visualization of algorithms through the use of high resolution color graphics.

Cornell University; Robert L. Constable; *Experiments With A Program Refinement System*; (DCR-8303327 A02); $63,000; 12 mos; (Joint support with the Software Engineering Program and the Intelligent Systems Program - Total Grant $189,494).

This research involves the judicious combination of the strong points of program development methodology and program verification in a logic for correct program development. This project is exploring a logic for formal top-down development called refinement logic. A system called Program Refinement Logic (PRL) has been designed and implemented. This system is being improved and expanded and considerable experimentation is being conducted.

Cornell University; David Gries and Fred B. Schneider; *Programming Methodology (Computer Research)*; (DCR-8320274 A01); $135,456; 12 mos.

This project continues research in five areas:

1. A theory of predicate transformers,
2. Synthesizing proof obligations for concurrent programs,
3. Investigation of trace logics,
4. A script-model editor, and
5. Sequential programming.

Brown University; Robert Sedgewick; *Analysis and Implementation of Algorithms (Computer Research)*; (DCR-8308806 A01); $55,833; 12 mos; (Joint support with the Theoretical Computer Science Program - Total Grant $111,666).

This research concerns the detailed study of the properties of fundamental algorithms. The research centers on drawing specific conclusions on the best way to solve important practical problems, backed up by careful implementations of basic algorithms, full mathematical analysis, and empirical studies to verify the analysis. At best this type of careful scrutiny often leads to the design of new and more efficient algorithms or to the identification of general techniques for use in studying algorithms. The research contributes to the understanding of important methods of problem solution.

Typically, research activity in this field spans a wide range of areas, from topics such as number theory and complex analysis to topics such as programming languages and computer architecture as they relate to a narrower range of algorithms and problem domains, which are studied in detail (always with an eye toward those techniques which are generally applicable). This type of detailed analysis is reserved for the most important problems. Specifically, the research is developing the mathematical analysis necessary to obtain detailed computational complexity results. This is done in the context of important sorting, searching, and computational geometry algorithms. The research also makes heavy use of tools that allow the visualization of algorithms through the use of high resolution color graphics.

Distributed Computing

University of Arizona; Gregory R. Andrews and Richard D. Schlicting; *Concurrent and Fault-Tolerant Programming (Computer Science)*; (DCR-8402090 A01); $99,901; 12 mos.

This project is concerned with the design of language notations for concurrent systems and the development of techniques for the construction of fault-tolerant distributed software. In particular, it is investigating two specific topics: the design and implementation of a family of languages for concurrent programming called Darwin, and the use of fail-stop processors as a basis for fault-tolerant programming.

The goal of the Darwin project is to develop a collection of languages for the different levels of systems programming. The languages contain a common core, and each family member contains mechanisms specifically suited for programming one of four levels: kernels and interrupt handlers, shared-memory systems, distributed systems, and fault-tolerant systems.

A fail-stop processor is a virtual processor that suffers only one type of failure: a "crash". This project is investigating the development of software for execution of (an approximation of) such a machine. The basis for this software is fault-tolerant actions, a language construct that allows recovery protocols to be bound to arbitrary sequences of statements.

The two projects are related in that the highest-level member of Darwin includes mechanisms for fault-tolerant programming such as fault-tolerant actions.

Yale University; Michael J. Fischer; *Theory of Algorithms and Distributed Systems (Computer Research)*; (DCR-8405478 A01 & A02); $184,934; 12 mos.

Theory is developed in the area of algorithms and distributed systems. The goals are to apply theoretical methods and tools to practically-motivated problems arising in applied areas of computer science, particularly in distributed computing. Topics studied include methods of achieving fault-tolerance in unreliable systems, a logic of knowledge for describing protocols, and other problems in the area of algorithms and efficient data structures.

Georgia Institute of Technology; Martin S. McKendry; *An Operating System Architecture for Reliable Distributed Computing (Computer Research)*; (DCR-8316590 A01); $78,326; 12 mos.

The physical attributes of distributed systems indicate that it should be possible to construct programs that tolerate failures better and run faster than they would on single machines. However, despite advances in hardware technology, especially in local area networks, this is not currently possible. Why?

This project is based on the belief that the primary reason is that the architectural foundations of centralized systems cannot be extended to the distributed environment without also carrying over the limitations of centralized systems. In particular, conventional architectures do not account for failures. Means of supporting distributed programs that provide mechanisms needed to manage failures have been studied. The techniques are under development in the Clouds group and are based on the notions of objects and atomic actions. Objects (instances of abstract data types) are a means of encapsulating data and guaranteeing data invariants. Thus, they provide a basis for localizing the effects of failures. Actions are a generalization of the notion of a transaction. They provide the basis for synchronization, and they delimit the scope of failures. Actions consist of operations on objects.

The techniques being used offer extremely high benefits. Suport for recovery is integrated into the lowest levels of the system. Thus, all system components can be designed to survive failures. Reliable computations can then be constructed to reconfigure automatically to accommodate failures. The facilities provided simplify both application systems and the operating system itself.

Massachusetts Institute of Technology; William E. Weihl; *The Use of Atomic Actions in Concurrent, Fault-Tolerant Systems (Computer Research)*; (DCR-8510014); $107,688; 24 mos.

The goal of this research is to develop systematic techniques for building concurrent, fault-tolerant systems. Topics of interest cover a broad range, including concurrency control and recovery protocols, program specification, and program structure. This research focuses on an approach based on the use of atomic actions and atomic data types. The approach has the usual modularity benefits of abstract data types, as well as the simplicity afforded by atomic actions, namely a simple semantics for failures and easy reasoning about interactions among concurrent activities. In addition, the approach can be used to build systems that permit high levels of concurrency.

Earlier research focused on precisely defining "atomicity," and on understanding how to specify atomic data types. The goal now is to understand better both how to implement atomic data types, and also the limitations of atomicity — for what kinds of applications are atomic actions an appropriate implementation structure? Implementation issues that are being studied include efficient synchronization and recovery protocols, how to structure the implementation of an atomic data type, linguistic support for implementing atomic data types, and debugging.

University of Washington; Lawrence Snyder; *Software Systems for the CHiP Computer (Computer Research)*; (DCR-8416878); $150,043; 12 mos.

There are no adequate programming facilities for nonshared memory parallel computers. The most common languages are simply sequential languages with several extensions for specifying concurrent activity, and the Poker Parallel Programming Environment, the most novel entry to date and the product of earlier Blue CHiP research, provides only a first step towards breaking from the sequential-language-extension approach. In this research two approaches to parallel programming are being studied concurrently. The first study seeks a method by which a shared memory programming model (e.g., a known programming language such as APL) can be implemented on a nonshared memory machine (e.g., the Configurable, Highly Parallel Computer). The second study seeks a new programming language, based on a nonshared memory model, that is expressive, convenient, parsimonious and efficient. Each language study utilizes theoretical and experimental techniques. Both efforts are influenced by a companion study of parallel algorithms and parallel programming paradigms. This research represents the theory, algorithms and software components of the ongoing Blue CHiP Project.

University of Wisconsin - Madison; Charles N. Fischer; *The Next Generation of Program Development Environments (Computer Research)*; (DCR-8502897); $48,255; 12 mos.

This project is investigating the next generation of language-based editors (LBEs) and program development environments (PDEs). Topics being studied include:

1. Automatic inference of the type of variables based on context;
2. Extension of PDE diagnostics to include language pragmatics; and
3. Interactive syntactic error repair.

Ways of extending the performance and capabilities of PDE are also being studied. Topics being studied include:

1. Automatic PDE interpreter generation from formal specifications; and
2. Priority-driven attribute evaluation for highly interactive semantic analysis.

University of Wisconsin - Madison; Mary Vernon; *A Graph Model Approach in Analyzing Performance and Deadlock In Distributed Systems (Computer Research)*; (DCR-8402680); $111,156; 24 mos.

This research seeks stronger methods for modeling and analysis of distributed systems. A long range goal is to define a model of distributed computation which permits direct analysis of stochastic behavior for performance estimation at the same time as it permits formal analysis of functional correctness. A corollary long range goal is to be able to incrementally add complexity to distributed systems models so that assumptions made at various levels of analysis and design are more easily validated. In the shorter term, recent work on the use of a formal graph model to analyze both performance and potential for deadlock in distributed computations is being extended. Primary attention is being given to applying queueing analysis techniques directly to the graph model, and to extending the reachability algorithms for detecting deadlock to include more system functionality. Work focuses in the following areas:

1. Direct application of queueing analysis methods to the graph model;
2. Extensions to the modeling primitives to include further features of queueing systems;
3. Application of the modeling, simulation, and analysis tools to experimental distributed algorithms which have the characteristic that both performance and deadlock are important issues;
4. Experience with the relationship between the structure of queueing network models and control flow models for deadlock detection; and
5. Introduction of formal language constructs to define data state changes in the graph model.

Verification

Artificial Intelligence Research Institute - Texas; Frank M. Brown; *Verification of Advanced Language (Computer Research)*; (DCR-8402412 A01); $52,579; 12 mos.

The overall goal of this project is an integrated deductive system applicable to program verification starting with an existing system called SYMEVAL. To reach this goal many difficult problems need to be solved, including the proof by induction of fully quantified theorems and the formalization of nondeterministic programming-language features.

BLANK PAGE

Special Projects

Privacy & Security

Stanford University; Martin E. Hellman; *Cryptography and Data Security (Computer Research)*; (ECS-8307741 A02); 12 mos; (Joint support with the Communication Systems & Signal Processing Program - Total Grant $74,509).

Because of the rapidly decreasing cost of computation, cryptography is important in the protection of valuable commercial and personal data. The research being pursued is motivated in part by previous work and in part by new directions to be investigated. The project involves the investigation of the following areas:

1. Time-memory processor tradeoffs,
2. Properties of random graphs,
3. Synchronous keystream systems,
4. Probabilistic approaches to cryptoanalysis,
5. Multiple encryption schemes, and
6. More efficient public key distribution systems.

Using previous results from time-memory-processor tradeoff, the very large scale integration (VLSI) packing problem - given a large number of circuit elements, arrange them in as small an area as possible - is being investigated. Further study of the hardware implementation of cryptoanalyis algorithms is expected to shed light on the distribution of drainage property of random graphs. A study to find the exact level of security of keystream systems based upon counters driving a general block cryptosystem is being undertaken.

University of California - Berkeley; Manuel Blum; *Transaction Protection Protocols (Computer Research)*; (DCR-8204506 A04); $93,585; 12 mos.

The objective of this research is to develop, study, and experiment with protocols for protecting electronic business transactions. A number of such protocols have been developed for sending certified mail, signing contracts, and exchanging secrets. A problem in each case is to achieve the equivalent of PARALLELISM through SEQUENTIAL exchange of messages. Signing contracts, for example, normally requires that the two parties to the contract be present at the same place and at the same time for signing. Signing contracts by telephone is difficult, even with computers, since whoever signs first obligates himself herself first. It is the task of the protocol to achieve the EQUIVALENT of simultaneous signing. This is possible using techniques of number theory and complexity theory that have been developed recently for work on public-key cryptography and signature schemes. The investigators are developing protocols for various problems such as these and are improving those already available in hopes of making them practical. With a few notable exceptions, the published literature does not give proofs of correctness and security, though these are starting to be developed. The investigators are contributing to the development of such proofs.

University of Wisconsin - Milwaukee; George I. Davida and Gilbert G. Walter; *Analog Data Protection (Computer Research)*; (DCR-8504620); $19,999; 12 mos; (Joint support with the Electrical and Optical Communications Program - Total Grant $39,999).

Cryptography has become an important tool in data protection. Cryptosystems can be used to solve a large number of the data security problems. While cryptosystems have been used successfully in protection of messages, the discovery of many different types of cryptosystems has lead to possible solutions for many other problems in data security that appeared to be either difficult to solve or were not suitable for "classical" types of cryptosystems. Most of the recent discoveries of cryptosystems have been in the digital domain. Digitizing data, however, is not always possible or economical. This research is intended to study new cryptosystems for protecting analog signals' privacy and authenticity. Public key analog cryptosystems are also being studied. In view of the relative scarcity of fast, secure digital public key cryptosystems, such cryptosystems might also be of interest in digital applications.

University of Wisconsin - Milwaukee; George I. Davida; *Data Security (Computer Research)*; (DCR-8112345 A05); $33,701; 6 mos.

The University of Wisconsin-Milwaukee is conducting research in the area of data security. The project concentrates on the the topics of operating systems security, encrypted relational databasess, and key management problems. The researchers are interested in designing security issues being addressed.

The Foundation provides opportunities for members of small-college science and engineering faculties to participate in research under the aegis of NSF investigators at research institutions. A faculty member in

the Department of Mathematics and Computer Science at Beloit College is participating in this research effort. Dr. Linda L. Deneen is collaborating with the principal investigator, Dr. George Davida, to develop new public-key encryption systems and to extend the research in providing effective system security.

The Societal Impact of Computing

University of Pittsburgh; Siegfried Treu; *Facilitative but Unconstraining User-Network Interface Systems in Office Environments*; (DCR-8505062); $54,855; 12 mos.

The focus of this applied research project is on the design and development of powerful interface systems to serve the users of multi-resource office networks. A testbed for such prototype systems is based on already acquired computer and networking hardware and software. Emphasis is being placed in part on finding and implementing means and methods for rendering the interface "facilitative", especially for the non-computer specialist working in the office. However, such facilitation is not to "constrain" the user into serving only simple-minded and routine computer-aided tasks; our technology should enable carefully staged advancement to more challenging and sophisticated tasks. Accordingly, this research focuses on the design of integrative interfaces that provide for certain kinds of supportive network context, adapt to the level of user advancement, and facilitate the interaction whenever appropriate.

Modelling & Simulation

University of Arizona; Bernard P. Zeigler; *Distributed Simulation of Hierarchical, Multicomponent Models (Computer Research)*; (DCR-8514348); $121,722; 12 mos.

Efforts are being continued to elaborate a methodology in which multicomponent discrete event models may be simulated employing multiprocessor architectures. A major part of the research concerns itself with the mapping of hierarchical multicomponent models onto distributed simulators so that correct and efficient simulation results. Text classes of models are being generated from case studies of computer networks as well as employing properties expressable within the discrete event formalism. Procedures for simplification of multilevel models so as to fit the constraints of multiprocessor architectures are also being developed. Theoretical concepts and results are cross-checked with experiments on models of distributed simulators expressed in conventional simulation languages and, when it becomes available, on a commercial network-in-a-box multiprocessor system.

The results of the research contribute to the computer science area by increasing the understanding of the utilization of a potentially powerful class of architectures and to the systems science area because of the extensions that are made to the theory of system representation so that it applies to the enlarged domain of distributed computer simulation.

Wayne State University; Bernard P. Zeigler; *Distributed Simulation of Hierarchical, Multicomponent Models (Computer Research)*; (DCR-8407230 A01); $45,440.

Efforts are being continued to elaborate a methodology in which multicomponent discrete event models may be simulated employing multiprocessor architectures. A major part of the research concerns itself with the mapping of hierarchical multicomponent models onto distributed simulators so that correct and efficient simulation results. Text classes of models are being generated from case studies of computer networks as well as employing properties expressable within the discrete event formalism. Procedures for simplification of multilevel models so as to fit the constraints of multiprocessor architectures are also being developed. Theoretical concepts and results are cross-checked with experiments on models of distributed simulators expressed in conventional simulation languages and, when it becomes available, on a commercial network-in-a-box multiprocessor system.

The results of the research contribute to the computer science area by increasing the understanding of the utilization of a potentially powerful class of architectures and to the systems science area because of the extensions that are made to the theory of system representation so that it applies to the enlarged domain of distributed computer simulation.

Networking, Computer Communications, and Distributed Computing

Purdue University; Douglas E. Comer, J. Tim Korb and Walter Tichy; *High-Level Network Protocols (Computer Research)*; (DCR-8219178 A02); $100,000; 12 mos; (Joint support with the Coordinated Experimental Research Program - Total Grant $275,358).

This research centers on three major problems occurring in distributed and computer networks. The first area of research is on algorithms and data structures to support electronic communication. As electronic mail services provide more information flow, ways to manage this information must be found. The second project concerns integrating the personal workstation into a general network where the computational resources of the entire network are available to the user. The final project deals with the problem of managing data concerning project development in a general distributed environment. As automation aids become more available, the activities of many users in a network which has both shared and independent computational and storage resources must be managed and coordinated.

University of New Orleans; Wayne Patterson; *RUI: CSNET Services (Computer Research)*; (DCR-8514025); $10,875; 36 mos.

As part of the Foundation's initiative to strengthen research environments in primarily undergraduate institutions, the Division of Computer Research considers proposals to defray CSNET dues and operating expenses under the research instrumentation category of the Research in Undergraduate Institutions (RUI) Program. The University of New Orleans has met both RUI and CSNET membership requirements and demonstrates a genuine research need in computer science. Approximately five faculty members at the university now have access to CSNET services. The benefits from network communication are significant and institutional commitment to future network use is evident.

The technical and administrative liaison for the CSNET site is Dr. Wayne Patterson, Chairman of the Department of Computer Science. Dr. Patterson and the four faculty members involved are highly competent and the research issues being addressed are important to the computer science area.

Brandeis University; James A. Storer; *Data Compression Algorithms Suitable for Communications and Parallel Implementations (Computer Research)*; (DCR-8403244); $36,110; 12 mos.

With many types of networks (e.g., distributed computing and electronic mail) communication channels are relatively slow. The ability to put large amounts of processing on a single chip promises to make sophisticated data compression algorithms truly practical. A data encoding/decoding chip can be placed at the ends of every communication channel, with no computational overhead incurred by the communicating processes. This research concerns the design of algorithms that can be used by such chips, as well as the parallel implementation of such algorithms.

Massachusette Institute of Technology; Robert G. Gallager and Pierre A. Humblet; *The Dynamics of Data Network Research (Computer Research)*; (ECS-8310698 A01); $33,500; 12 mos; (Joint support with the Communication Systems and Signal Processing Program - Total Grant $126,000).

This group has approached their research on communication networks by developing, in a fundamental way, new models and new algorithmic and analytic techniques. This study involves numerous problems the solution of which will reduce the ad-hoc procedures employed in the design of networks. Their approach requires that each individual investigator have an understanding of communication and information theory, computer science, control theory, and operations research. Their trademark is a thorough, indepth investigation which results in a fundamental contribution. As a consequence of this approach, their research commitments are long-range.

University of Massachusetts - Amherst; Victor R. Lesser; *Coordination in Cooperative Distributed Problem Solving Systems (Computer Research)*; (DCR-8300239 A02); $22,000; 12 mos; (Joint support with the Computer Systems Design Program and the Intelligent Systems Program - Total Grant $128,696).

A key problem in cooperative distributed problem solving is developing network coordination policies that provide sufficient global coherence for effective cooperation. This problem is difficult because limited internode communication precludes the use of a global "controller" node. Instead, each node must be able to direct its own activities in concert with other nodes, based on incomplete, inaccurate and inconsistent information. This requires a node to make sophisticated local decisions that balance its own perceptions of appropriate problem solving activity with activities deemed important by other nodes. Current research at the University of Massachusetts has included the development of a node architecture capable of such sophisticated local decision making. This architecture has been implemented as part of the distributed vehicle monitoring testbed: a flexible and fully instrumented research tool for the empirical evaluation of distributed network designs and coordination policies. This re-

search builds on this node architecture and testbed to explore (through actual implementation and empirical studies) a variety of approaches to network coordination that include organizational self-design, distributed load-balancing, negotiation among nodes, planning of internode communication strategies, and knowledge-based fault-tolerance.

University of Massachusetts - Amherst; Edward M. Riseman and Arthur S. Gaylord; *A Group Research Facility for Artificial Intelligence, Distributed Computing, and Software Systems (Computer Research)*; (DCR-8318776 A02); $70,000; 12 mos; (Joint support with the Intelligent Systems Program, and the Software Engineering Program - Total Grant $120,000).

Experimental computer research is focused in the six areas of computer vision, robotics, distributed computing, software development, intelligent interfaces, and natural language processing. While most of the actual research activities are supported under other grants and contracts, including some from NSF, this grant will provide partial support for the experimental research facility used by all the researchers, supplementing institutional, departmental, and other funds. NSF support is used for administration and operation of the COINS research facility, with the funding going towards technical support staff and equipment maintenance.

University of Massachusetts - Boston; Elizabeth J. O'Neil; *RUI: CSNET Services (Computer Research)*; (DCR-8511438); $8,200; 36 mos.

As part of the Foundation's initiative to strengthen research environments in primarily undergraduate institutions, the Division of Computer Research considers proposals to defray CSNET dues and operating expenses under the research instrumentation category of the Research in Undergraduate Institutions (RUI) Program. The University of Massachusetts-Boston has met both RUI and CSNET membership requirements and demonstrates a genuine research need in computer science. Approximately seven faculty members at the university have access to CSNET services. The benefits from network communication are significant and institutional commitment to future network use is evident. The technical and administrative liaison for the CSNET site is Dr. Elizabeth O'Neil, Professor of Mathematics and Computer Science.

Michigan State University; Herman D. Hughes; *Analysis of Priority Schemes for Local Networks (Computer Research)*; (DCR-8318529 A01); $61,885; 12 mos.

The advent of large scale integration (LSI) and very large scale integration (VLSI) devices provides inexpensive high-bandwidth communication between computers within a limited geographic territory, and is a significant event in the history of computing. The new technology has given rise to a proliferation of local area networks (LANs) during recent years, and this trend seems likely to continue during the next decade. As LANs and their applications (e.g. office automation, real-time processes, research and development, educational activities, etc.) increase in the United States and elsewhere, it becomes very important that they are designed to meet certain performance goals for a wide range of services (voice, video, facsimile, etc.). Three important design parameters for LANs are:

1. The transmission media,
2. The topologies, and
3. The channel access protocols.

These parameters determine whether or not a given LAN can achieve certain performance goals.

This study is examining the design parameters for LANs recommended by the IEEE Standards Committee in regard to their performance (i.e., channel throughput, queueing delays) for integrated services, and publish reports on how these parameters may be altered (e.g., changing baseband to broadband, assigning priority to packets) to allow a single communication network to yield acceptable performances.

Michigan Technological University; David A. Poplawski; *RUI: CSNET Services (Computer Research)*; (DCR-8511439); $8,000; 36 mos.

As part of the Foundation's initiative to strengthen research environments in primarily undergraduate institutions, the Division of Computer Research considers proposals to defray CSNET dues and operating expenses under the research instrumentation category of the Research in Undergraduate Institutions (RUI) Program. Michigan Technological University has met both RUI and CSNET membership requirements and demonstrates a genuine research need in computer science. Approximately six faculty members at the University have access to CSNET services. The benefits from network communication are significant and institutional commitment to future network use is evident. The technical and administrative liaison for the CSNET site is Dr. David Poplawski, Professor of Mathematics and Computer Science.

Carleton College; Richard Nau, Richard Brown, Kathleen Galotti, Roger Kirchner, and Linda Ness; *RUI: CSNET Services (Computer Research)*; (DCR-8503284); $6,562; 36 mos.

As part of the Foundation's initiative to strengthen research environments in primarily undergraduate institutions, the Division of Computer Research considers proposals to defray CSNET dues and operating expenses under the Research instrumentation category of the Research in Undergraduate Institutions (RUI) Program. Carleton College has met both RUI and CSNET membership requirements and demonstrates a

genuine research need in computer science. Approximately five faculty members at the College have access to CSNET services. The benefits from network communication are significant and institutional commitment to future network use is evident.

The technical and administrative liaison for the CSNET site is Dr. Richard W. Nau, Professor of Mathematics and Computer Science. Dr. Nau and the four faculty members involved are highly competent and the research issues being addressed are important to the computer science area.

Wake Forest University; David J. John; *RUI: CSNET Services (Computer Research)*; (DCR-8506092); $8,850; 36 mos.

As part of the Foundation's initiative to strengthen research environments in primarily undergraduate institutions, the Division of Computer Research considers proposals to defray CSNET dues and operating expenses under the research instrumentation category of the Research in Undergraduate Institutions (RUI) Program. Wake Forest University has met both RUI and CSNET membership requirements and demonstrates a genuine research need in computer science. Approximately four faculty members at the University have access to CSNET services. The benefits from network communication are significant and institutional commitment to future network use is evident. The technical and administrative liaison for the CSNET site is Dr. David J. John, Professor of Mathematics and Computer Science.

Canisius College; James A. Leone; *RUI: CSNET Services (Computer Research)*; (DCR-8511659); $7,800; 36 mos.

As part of the Foundation's initiative to strengthen research environments in primarily undergraduate institutions, the Division of Computer Research considers proposals to defray CSNET dues and operating expenses under the research instrumentation category of the Research in Undergraduate Institutions (RUI) Program. Canisius College has met both RUI and CSNET membership requirements and demonstrates a genuine research need in computer science. Approximately four faculty members at the College have access to CSNET services. The benefits from network communication are significant and institutional commitment to future network use is evident. The technical and administrative liaison for the CSNET site is Dr. James Leone, Professor of Computer Science. Dr. Leone and the three faculty members involved are highly competent and the research issues being addressed are important to the computer science area.

New York University; Zvi M. Kedem; *Theory and Applications of Database Concurrency Control (Computer Research)*; (DCR-8416422); $7,465; 3 mos.

Recent research has led to the development of non-two-phase locking protocols for controlling concurrency which assure correct results in a database without causing deadlocks. Such systems make no use of information about the operations performed by transactions on data entities, although they do use apriori information about the order in which transactions may access the entities. If information of the former type is also used, concurrency may be increased, but deadlocks are more difficult to avoid. This research is investigating non-two-phase locking protocols for such 'partially interpreted' systems with the goal of using both semantic and syntactic information to increase concurrency and possibly also avoid deadlocks. Mechanisms for removing deadlocks when avoidance is impossible or too costly are being studied. Methods for conditioning the structure of a database to assure proper accessing behavior of the transactions so that such protocols can be used, as well as new general means for allowing concurrent dynamic modifications of the structure of the database, are being sought. The performance of all mechanisms developed is being evaluated in order to ascertain their applicability to real database systems.

University of Texas - Austin; Simon Lam; *Protocols for Communication and Resource Allocation (Computer Research)*; (ECS-8304734 A02); $20,000; 12 mos; (Joint support with the Communication & Signal Processing Program - Total Grant $78,715).

Resource allocation techniques and their implementation in a distributed network environment are being investigated. Networks based upon virtual channels are currently the dominant packet networks. A new model is used to study the performance of these networks. Specifically, problems of congestion control, routing, aid-to-aid flow control, and the effect of their interaction on network performance are being examined. Techniques for synchronizing network status information in both broadcast and non-broadcast networks, and scheduling in broadcast networks with limited synchronization are being sought. A hierarchial model is being used to develop methods to analyze communication protocols.

Western Washington University; James L. Johnson; *RUI: CSNET Services (Computer Research)*; (DCR-8512778); $13,750; 36 mos.

As part of the Foundation's initiative to strengthen research environments in primarily undergraduate institutions, the Division of Computer Research considers proposals to defray CSNET dues and operating exppenses under the research instrumentation category of the Research in Undergraduate Institutions (RUI) Program. Western Washington University has met both RUI and CSNET membership requirements and dem-

onstrates a genuine research need in computer science. Approximately five faculty members at the University have access to CSNET services. The benefits from network communication are significant and institutional commitment to future network use is evident.

The technical and administrative liaison for the CSNET site is Dr. James L. Johnson, Professor of Computer Science. Dr. Johnson and the four faculty members involved are highly competent and the research issues being addressed are important to the computer science area.

Other Projects, Symposia, Colloquia, and Studies

Stanford University; Stanley Peters; *Workshop on Information and Representation, Washington, D.C., March 31 - April 1, 1985*; (BNS-8512693); $2,927; 3 mos; (Joint support with the Linguistics Program - Total Grant $14,634).

It has been clear for some time now to observers of the scientific scene that there is a great deal of intellectual ferment at the confluence of several disciplines: cognitive psychology, information science, linguistics, computer science, analytic philosophy, and artificial intelligence. Workers in each of these areas are discovering relevant results and methods in the others, and numerous collaborative and interdisciplinary enterprises are being undertaken.

This two-day workshop examined this emerging area of activity to identify its focus and its boundaries and to explore what resources are presently available for research in the area and what may be needed in the future if it is to meet its potential. The participants in the workshop are all leaders in the several research disciplines involved and all are active in research of the interdisciplinary sort which was examined.

University of California, San Diego; John F. Asmus; *Computerized Digital Image Enhancement Applied to Paintings (Computer Research)*; (DCR-8415095); $35,370; 12 mos.

Frequently, the features of fine paintings have faded or are obscured by dirt, darkened varnish, craquelure, foxing, and earlier restorations. Scholarly study and preconservation inspections are impeded by such blemishes in many instances. This enterprise seeks to identify those digital image processing techniques developed over the past decade that may be applied to the alleviation of conservation and interpretation problems of important pictures. Working with colleagues at museums together with associates at leading image processing laboratories, the research team is surveying existing computer software and expertise in order to match museum image problems with potential state-of-the-art solutions. Case studies are being selected and trial experiments are being performed to aid the conservation/interpretation efforts, reveal the potential of the approach, and expose the commmunity to the full power of this opportunity.

University of Chicago; Robert Soare and Laszlo Babai; *Workshop on Computational Complexity, May 2-4, 1985, University of Chicago, Chicago, Illinois (Computer Research and Mathematical Sciences)*; (DCR-8507772); $10,082; 3 mos; (Joint support with the Mathematics Special Projects Program - Total Grant $12,082).

The Department of Computer Science of the University of Chicago held a Workshop on Computational Complexity on May 2-4, 1985. The workshop brought together a number of researchers to discuss topics of current interest in central areas of Computational Complexity Theory. The core of the meeting was a series of four expository lectures by Richard M. Karp on parallel computation, one of the fastest growing areas of the theory of computing. In addition, four one-hour lectures were given by R. L. Graham, F. T. Leighton, L. Lovasz and E. Szemeredi. The unusual concentration of research potential in the Theory of Computing represented by the permanent and long-term visiting members (including Lovasz, Szemeredi, Babai, A. Lenstra, G. Schnitger, A. Itai) of the Department made the University of Chicago a particularly suitable place for the meeting.

University of Kentucky; Wiktor Marek; *Workshop in Logic and Computer Science, Lexington, Kentucky, December, 1984 (Mathematical Sciences and Computer Research)*; (DMS-8403701); $10,000; 6 mos; (Joint support with the Mathematics Special Projects Program - Total Grant $14,500).

This workshop was devoted to applications of logic in computer science, in particular, applications to the theory of databases and artificial intelligence. The principal topics discussed were logic of databases, changing database systems, data structures for databases, query languages, artificial intelligence considerations, computational aspects of databases, complexity of database problems, and database-oriented formal languages.

Massachusetts Institute of Technology; Silvio Micali; *A Workshop on the Mathematical Theory of Security, Boston, Massachusetts, June 2-5, 1985 (Computer Research)*; (DCR-8506174); $35,000; 3 mos.

The aim of this workshop was to coordinate basic research in the mathematical theory of secure communications and protocols. Specific goals were to:

1. Provide the necessary mathematical framework to discuss new security issues,
2. Set desired standards and offer helpful guidelines by selecting archetypical problems and providing concrete solutions, and
3. Generate momentum and critical mass and broaden the communication channel between the research community of the mathematical theory of security and the research communities of the related fields - theory of computation, logic, information theory, and number theory.

University of Massachusetts - Boston; Robert Morris; *RUI: Studies in Digital Typography (Computer Research)*; (DCR -8503154); $60,547; 24 mos.

This research involves two areas of study. The first is concerned with problems arising in matching digital fonts of two different resolutions, e.g., a CRT font with a printer font. Inadequately matched setwidths result in radically different appearances of the same lines on the two devices. Such problems reduce the utility of "what-you-see-is-what-you-get" editors, composition preview systems, or laser printer proofing of typesetter composition. The investigators are constructing and evaluating mismatch measures, and seeking ways to incorporate them into composition systems. The second area studies the digitization of fonts from outlines. Others will study the utility of conic versus cubic splines, some problems in aliasing, and some digitization algorithms.

Clarkson College of Technology; Edward Dubinsky; *Institute for Retraining in Computer Science*; (DMS-8318765); $2,500; 12 mos; (Joint support with the Mathematics Special Projects Program - Total Grant $5,000).

There is a national shortage of college faculty in computer science that has reached alarming proportions. Although undergraduate student demand for computer science courses is growing rapidly, the number of new Ph.D.'s in this area that are taking academic positions each year has not increased since 1975.

The Institute for Retraining in Computer Sciences is based on the idea that college faculty in other areas where the personnel situation is not so critical can make use of their background in subjects such as mathematics, their teaching experience and what knowledge they do have of computer science to learn quickly enough material so as to be qualified to teach a significant portion of the undergraduate curriculum in computer science.

A high quality program using outstanding computer scientists from major departments as faculty has been designed specifically for this institute and has been implemented with a pilot group.

University of Oklahoma; Sudarshan K. Dhall and S. Lakshmivarahan; *A Workshop on Parallel Processing Using the HEP, March 20-21, 1985, University of Oklahoma, Norman, Oklahoma (Computer Research)*; (DCR-8500481); $2,984; 6 mos; (Joint support with the Office of Advanced Scientific Computing - Total Grant $4,984).

The University of Oklahoma conducted this first workshop on "Parallel Processing Using the Denelcor Heterogeneous Element Processor (HEP)." The workshop was held March 20-21, 1985, at the University of Oklahoma in Norman, Oklahoma. The aim of the workshop was to bring together:

1. The group of hardware and software designers of the Heterogeneous Element Processor;
2. Researchers from various universities and laboratories who are currently using and engaged in further exploring the full potential of the HEP in large-scale scientific and engineering problem solving;
3. Representatives from various federal agencies engaged in fostering research, education, and training in large-scale computing; and
4. Researchers from various industries with a view to evaluate the present state-of-the-art and future promise for supercomputing using the HEP.

The topics covered ranged from the architecture and software of the HEP to applications in Artificial Intelligence and Expert Systems. Contributed and expository research papers are being published in the form of a Proceedings.

American Federation of Information Processing Societies, Inc.; Richard H. Austing; *Group Travel for U.S. Participants in the 1986 Congress of the International Federation for Information Processing, September 1-5, 1986, Dublin, Ireland (Computer Research)*; (DCR-8504191); $40,000; 18 mos; (Joint support with the Division of Information Science and Technology - Total Grant $60,000).

The tenth triennial Congress of the International Federation for Information Processing (IFIP) will be held September 1-5, 1986 in Dublin, Ireland. The IFIP Congresses have provided an international forum in which ideas could be exchanged across areas of the disciplines and, as a result, they have been recognized as the most prestigious of the international meetings in the information science field. The Congresses provide an exceptional opportunity for exchange of information at the frontiers of research between the most productive specialists. United States scientists become acquainted with the current research efforts of their colleagues in all countries. This immediate acquisition of information overcomes the substantial delay in the appearance of new results in the professional literature.

The American Federation of Information Processing Societies (AFIPS) represents the information science interest in the United States. It is administering a travel grant program to support the travel of 75 information scientists from the U.S to the Congress. Their attendance should be of genuine benefit to their research potential which will contribute to the advancement of computer and computer-related research in the United States.

The project is under the direction of Dr. Richard Austing who has considerable experience in administering group travel grants of this nature.

University of Washington; Paul R. Young; *Growth in Academic Computer Science: An Impending Crisis?*; (DCR-8517081); $48,000; 12 mos.

The Computer Science Board is an independent organization consisting of 24 members representing Ph. D. granting Computer Science Departments and major Computer Science Research organizations in the United States. Four members of the Board (David Gries - Cornell University, Ray Miller - Georgia Tech, Robert Ritchie - Xerox PARC, and Paul Young - University of Washington and Committee Chair) have conducted a preliminary study of funding patterns for basic research inn computer science and presented these findings to the Computer Science Board in a report entitled "Growth in Academic Computer Science: An Impending Crisis?" Under the leadership of Dr. Young, the study is being extended:

1. To measure the impact of industrial funding;
2. By performing a more careful assessment than given in the previous study of the impact of CER funding;
3. By broadening the study to include computer engineering; and
4. By updating all study analyses with available new or recent year data.

Theoretical Computer Science

Concrete Complexity and Analysis of Algorithms

University of Alabama - Huntsville; Peter J. Slater; *NSF-CBMS Regional Conference on Applications of Combinatorics & Graph Theory to Computer Science, Huntsville, Alabama, June 1985 (Mathematical Sciences and Computer Research)*; (DMS-8503703); $1,000; 6 mos; (Joint support with the Mathematics Special Projects Program - Total Grant $23,900).

This proposal is for one of a number of regional conferences recommended for funding by the NSF, with supporting services for the processing of the proposals and for the conduct of the conferences being provided by the Conference Board on the Mathematical Sciences (CBMS) through the mechanism of a contract with NSF.

Criteria for the recommendation of funding of each proposal were:

1. Principal Speaker. The principal speaker is a world leader in current research in the topic of the conference, as well as an effective and able lecturer and expositor.
2. Subject Matter. The topic of the conference is of real significance for mathematics, and is one in which there exists considerable research activity.
3. Regionality. The topic of a conference was chosen from a field in which there is a need for stimulation and expansion of research activity in a broad geographic area (perhaps 300-500 miles in radius) surrounding the host institution.
4. Host Institution. It was deemed desirable for the host institution to have a modicum of competence in the subject area of a proposed conference in order that substantial benefits may accrue to the host institution's faculty. The principal speaker is not a member of the faculty of the host institution, nor is he (or she) from a neighboring institution.

Mathematical Sciences Research Institute, Inc.; Richard Karp and Stephen Smale; *Computational Complexity (Computer Research)*; (DCR-8418523); $60,000; 24 mos.

As part of the year-long program in Computational Complexity at the Mathematical Sciences Research Institute in 1985-86, the Institute will award competitively eight postdoctoral fellowships. The standard Institute stipend of $25,500, judged to be competitive in mathematics, is not competitive for theoretical computer scientists. The $60,000 awarded will supplement each of these eight fellowships in Complexity theory by an average of $7,500 to bring them up to the more competitive figure of $33,000. This is essential in order to attract the best young scholars in the field to participate in this timely and exciting program.

Stanford University; Andrew C. Yao; *Studies in Algorithms and Computational Complexity (Computer Research)*; (DCR-8308109 A03); $99,413; 12 mos.

Research is conducted along two general directions: the development and analysis of complexity models for classes of computational problems, and the probabilistic analysis of algorithms. Three objectives are emphasized:

1. The development of general unifying models which permit the study of many individual problems within the same framework;
2. The proof of upper and lower bounds for individual problems;
3. The development of novel mathematical techniques for solving problems.

Specific topics investigated include the complexity of data structure problems, generating nonuniform random numbers, distributed computing, VLSI computations, probabilistic behavior of algorithms, and probabilistic complexity.

University of California - Berkeley; Manuel Blum; *Cryptography and Computational Complexity (Computer Research)*; (DCR-8513926); $89,387; 12 mos.

This research focusses on the following group of topics involving a breadth of techniques:

1. The development of a method of program fingerprinting that can potentially be taught and used in a law court setting;
2. The pinpointing of specific sources of hardness in problems (This will include the study of carries as a source of hardness in numerical computations);
3. Generation of random numbers with randomness being assured through the use of large memory resources; and
4. The following relative of the P = NP problem; Is $P^A = NP^A \cap \text{Co-}NP^A$ with probability 1 for random function oracles A?

University of California - Berkeley; Michael Harrison; *Theory of Computing (Computer Research)*; (DCR-8311787 A01); $99,791; 12 mos.

Several research directions are pursued. These include operating system protection, formal language theory questions, and string rewriting systems.

In operating system protection, a general model for a protection system is formulated so that "safety" questions are mathematically precise and so that security policies can be formulated and tested. Additionally, the addition of cryptographic techniques to operating systems are being investigated.

In formal language theory and string rewriting systems questions related to generalized pumping lemmas, grammatical transformations, the complexity of parsing and recognition, and equivalence questions for DOL systems are studied.

University of California - Berkeley; Richard M. Karp; *The Design and Analysis of Combinatorial Algorithms (Computer Research)*; (DCR-8411954); $95,000; 12 mos.

This research concentrates on six major interrelated themes: probabilistic analysis of combinatorial algorithms, randomized algorithms, parallel algorithms, sorting and searching, approximation algorithms and combinatorial problems in VLSI design. Within these areas a large number of concrete research problems are studied.

University of California - Berkeley; Eugene L. Lawler and Charles U. Martel; *Combinatorial Algorithms (Computer Research)*; (DCR-8311422 A01); $103,272; 12 mos.

Efficient algorithms for sequencing and scheduling are studied in the general context of combinatorial optimization. A variety of important problems are studied including machine shop scheduling (important for the efficient utilization of scarce resources), the traveling salesman problem (important, for example, for the siting of refineries), and shortest path problems for large networks (important for computer communication problems). The goal is to develop fast algorithms for special cases of the above problems. Additionally, the theoretical framework for the generalized study of such problems is being investigated.

University of California - Irvine; George S. Lueker; *Probabilistic Methods in Computer Science (Computer Research)*; (DCR-8404898 A01); $112,870; 24 mos.

Professor Lueker is investigatinge problems using probabilistic methods. In particular, he is using probabilistic analysis of algorithms to:

1. Analyze problems that can be viewed as a selection of an optimal subgraph in a weighted graph (e.g., determining differences between independent and dependent weighting schemes);
2. Develop sorting algorithms that work in linear time under pathological input distributions;
3. Analyze packing problems, such as the Knapsack problem, bin-packing, and partitioning;
4. Study difficult classical problems, such as the analysis of hashing and dynamic storage allocation, in which recent advances offer some hope of significant progress; and
5. Study distributed processing, an area of current interest in which probabilistic methods are useful.

As a byproduct of such probabilistic analyses, a smaller component of the research is to devise algorithms that work well even in the worst case.

University of California - San Diego; Michael Fredman; *Studies in Data Structures and Related Algorithms (Computer Research)*; (DCR-8504245); $43,709; 12 mos.

Professor Fredman is developing methods in the data structure field that:

1. Apply refined complexity analysis to develop amortized complexity, self-adjustment, application-induced refinement, and inherent lower bounds (information content);
2. Construct efficient data structures with one (or few) bit changes and with constraints on the update task and that exploit special-purpose data structure hardware; and
3. Solve problems involving nearest neighbor queries in domains involving complicated distance functions.

University of Southern California; Leonard Adleman; *Computation Complexity and Its Relationship to Number Theory (Computer Research)*; (DCR-8022533 A04); $41,429; 12 mos.

Research is being conducted in three areas: computational complexity, cryptography, and VLSI.

1. Computational Complexity. Several aspects of computational complexity are being investigated with particular emphasis on its relationship to number theory. Much centers on the P = NP problem.
2. Cryptography. Public security and privacy.
3. VLSI. The VLSI related problems are: layout, design aids, design paradigms, and complexity.

University of Southern California; Gary L. Miller; *Algebraic Techniques in the Design of Efficient Algorithms (Computer Research)*; (DCR-8514961); $28,855; 6 mos.

This research focuses on the computational complexity of certain problems which have a strong algebraic component. Special attention is being given to certain enumerative problems including:

1. Extending previous work on solvability by radicals to questions in algebraic simplification; and
2. Counting the number of integer points in a polytope of bounded dimension.

University of Colorado - Boulder; Karl Winklmann; *Information-Flow Characteristics of Combinatorial Problems (Computer Research)*; (DCR-8500741); $3,304; 12 mos.

An interesting property of distributed computation is that only local information may be utilized during the computation. Non-deterministic polynomial-time bounded computations may share this property. Professor Winklmann is investigating this hypothesis. A notion of algorithmic information flow applicable to nondeterministic computations is being explored. NP-complete problems are being classified according to the complexity of their information flow. The long open P ≠ NP conjecture is being explored.

University of Georgia; Carl Pomerance, Jeffrey W. Smith, and Randy Tuler; *High Speed Factoring with the Quadratic Sieve Algorithm and a Pipeline Architecture (Computer Research and Mathematical Sciences)*; (DCR-8421341); $36,000; 12 mos; (Joint support with the Algebra and Number Theory Program and the National Security Agency - Total Grant $99,985).

Professors Pomerance, Smith and Tuler are simulating, designing and fabricating a special purpose processor to factor 90-digit numbers in less than one week and use this processor to measure the performance of the quadratic sieve algorithm, one of the best methods for factoring.

It is anticipated that such a factoring process can be used for:

1. Determining the security of public key cryptography schemas (for the National Security Agency);
2. Answering general problems dealing with complexity issues (for the theoretical computer science community); and
3. Developing new number theory techniques (for the algebra/number theory community).

Northwestern University; Der-Tsai Lee; *Geometric Complexity and Related Problems (Computer Research)*; (DCR-8420814); $36,497; 12 mos.

Professor Lee is:

1. Cataloging geometric problems that require (nlogn) time in an algebraic computation tree model;
2. Determining nontrivial lower bounds, other than (nlogn);
3. Considering the complexity of triangulating a simple polygon, determining the closest pair of vertices of a simple polygon and simplicity testing; and
4. Studying problems in the framework of constrained proximity, including geometric constructions, constrained shortest path problem, generalized Delaunay triangulation of planar straight-line graph, complexity of rectilinear art gallery problem, signature problem of rectilinear polygons and dynamic computational geometry.

University of Illinois - Chicago; Wolfgang Maass; *Lower Bound Arguments in Machine Based Complexity Theory (Computer Research)*; (DCR-8504247); $69,083; 24 mos.

Professor Maass is focusing on the development of new lower bound arguments in several areas of machine-based complexity theory. In particular, the applicability of powerful tools from combinatorics and mathematical logic is being analyzed. More specifically, the following problems are being investigated:

1. Stepwise development of lower bound arguments for two-tape Turing machines,
2. Exploration of new types of lower bound arguments for random access machines, and
3. Application of recently developed lower bound techniques for random access machines to models of highly parallel computers.

A likely outcome of this research is that the following question will be answered: Is the state-of-the-art data manipulation technique as much as can be done toward finding lower bounds, or are innovative new algorithms possible? Either outcome would be a significant advance in the understanding of how to proceed with complexity analyses.

University of Illinois - Urbana; C. L. Liu; *Combinatorial Algorithms (Computer Research)*; (DCR-8404239); $52,568; 12 mos.

This project is studying combinatorial algorithms under three categories: combinatorial optimization, optimization problems in design of integrated circuits, and combinatorics. In combinatorial optimization the problems are of the scheduling variety including an optimal time slot assignment problem, a personnel assignment problem, and a problem of initiation of maintenance cycles. In the design of integrated circuits, the optimal folding of programmable logic arrays is being considered as will the channel routing problem in which the routing algorithm is given a certain amount of freedom in selecting some of the pins to be connected.

Purdue University; Greg N. Frederickson; *Graph Algorithms and Data Structures (Computer Research)*; (DCR-8320124 A01); $46,378; 12 mos.

This research focuses on three interrelated areas: data structures for update problems, algorithms for graph problems, and complexity issues in distributed

computing. The common theme is the investigation of efficient computation in these areas. The generation of improved proof techniques and the characterization of significant complexity relationships are the goals of this research.

Johns Hopkins University; S. Rao Kosaraju; *Applications of Foundations of Computing (Computer Research)*; (DCR-8506361); $47,759; 12 mos.

Professor Kosaraju is investigating problems in three areas of computer science:

1. VLSI problems, including trade-offs with and without constraints on flow information and compaction of certain types of VLSI layout on meshes;
2. Parallel algorithm problems, including examination of natural ways of designing algorithms that can be automatically translated into parallel algorithms (e.g., to produce systolic circuits); and
3. Motion planning problems, including investigation of regions of the plane that can be reached by a 2-dimensional set of linkages under various constraints.

In the area of robotics, it is anticipated that the theory of nondeterministic automata will play an important role in the structural analysis of reachable regions of a robot arm.

Massachusetts Institute of Technology; Shafrira Goldwasser; *Computational Complexity Based Cryptography (Computer Research)*; (DCR-8509905); $103,435; 24 mos.

Professor Goldwasser is developing a mathematical theory of cryptography and is exploring its relations to reliable distributed computing, information theory, and randomness. More specifically, she is:

1. Strengthening basic cryptographic tools by extending their security against more powerful and realistic adversaries than previously considered (important to commerce, government, defense, etc.);
2. Designing fault-tolerant distributed algorithms (important for parallel computing);
3. Providing a computational complexity measure of knowledge and using it to prove correctness of cryptographic protocols; and
4. Exploring applications of pseudo-random functions to complexity theory.

Massachusetts Institute of Technology; Silvio Micali; *Randomness and Computational Difficulty: Theory and Applications (Computer Research)*; (DCR-8413577); $49,031; 12 mos.

Randomness and computation play dual roles and interact in many ways. A coherent theory of randomness based on computational difficulty is being developed. Particular attention is devoted to finding the most powerful properties of randomness that can be achieved by polynomial-time simulations. This randomness theory is being applied to cryptography and cryptographic protocols, complexity theory and distributed computing.

Massachusetts Institute of Technology; Gary L. Miller; *Algebraic Techniques in the Design of Efficient Algorithms (Computer Research)*; (DCR-8404188); $45,303; 12 mos.

This research focuses on the computational complexity of certain problems which have a strong algebraic component. Special attention is being given to certain enumerative problems including:

1. Extending previous work on solvability by radicals to questions in algebraic simplification; and
2. Counting the number of integer points in a polytope of bounded dimension.

Massachusetts Institute of Technology; Ronald L. Rivest; *Concrete Computational Complexity (Computer Research)*; (DCR-8006938 A05); $96,998; 12 mos.

This research is concerned with the study of algorithms and computational complexity in the area of VLSI design and the theory of cryptology. The approach to VLSI design involves developing component packing algorithms and wire routing algorithms subject to the constraints imposed by VLSI manufacturing techniques. The approach to cryptology research involves the development of a conceptual framework that allows the measurement of the security of cryptosystems. This conceptual framework rests upon the development of a theory of "finite" computational complexity.

Massachusetts Institute of Technology; Michael Sipser; *Computational Complexity and Algorithms (Computer Research)*; (DCR-8304769 A01); $37,972; 12 mos.

The primary focus of this research project is the study of the inherent complexity of problems. Of particular interest are problems which have apparent exponential complexity yet for which no superpolynomial lower bound is known. The emphasis is on the development of techniques which may give insight into fundamental questions about polynomial versus exponential growth such as the P = NP question.

The area of probabilistic computation is also being explored. A useful definition of the notion "random sequence" via complexity theory is sought. Applications and properties of such sequences are being investigated.

University of Minnesota; Oscar H. Ibarra; *Problems in Computational Complexity (Computer Research)*; (DCR-8304756 A03); $64,480; 12 mos.

Research is being conducted in the following areas:

1. Characterizations and computational complexity of some models for VLSI;
2. Automata-based complexity theory: complexity measure trade-offs, nondeterminism versus determinism, hierarchies of computation, characterizations of complexity classes; and
3. Decision problems in automata theory, formal languages, and other areas of computer science.

University of Minnesota; Sartaj K. Sahni; *Algorithmic Study of Combinatorial Problems (Computer Research)*; (DCR-8305567 A03); 10 hours of supercomputer time; 12 mos.

Combinatorial problems are problems that deal with collections of objects in order to obtain information about the structures of those collections. In theoretical computer science the combinatorial problems are abstracted from problems that are interesting in the sense that they solve real world problems. An example of such a problem is a wire routing problem for VLSI circuits where it is desirable to connect circuit elements together minimizing the longest wire length between circuit elements. Typically combinatorial problems are computationally difficult to solve exactly and thus approximate solutions are sought.

University of Nebraska - Lincoln; Sharad C. Seth; *Decoding Substitution Ciphers Using a Spelling Dictionary with an Application to Unsupervised Character Classification (Computer Research and Information Science)*; (DCR-8421162); $26,943; 12 mos; (Joint support with the Division of Information Science and Technology - Total Grant $53,885).

Professors Seth and Nagy are studying the decoding of substitution ciphers (using a spelling dictionary) and are applying them to unsupervised character classification. This work has important applications in optical character recognition.

The formulation of the computer science problem includes investigation of search algorithms, linguistic heuristics, dictionary organization, and string matching. The theoretical problem that is being explored is the relationship between the length of the dictionary, length of cipher text, and the probability of correct assignment.

The advantages of such a dictionary-based approach lie in its potential for solving the translation (of text to standard code) using relatively small text samples and spelling dictionaries. This approach is markedly different from such commercial products as the Kurzweil Data Entry Machine. A likely outcome of this research is that it will make optical character recognition more powerful, more flexible, and much cheaper. The implications for aids to the blind are highly significant, to mention only one way in which optical character recognition impacts society.

Columbia University; Zvi Galil; *The Design of Efficient Algorithms and Proving Lower Bounds (Computer Research)*; (DCR-8303139 A02); $57,455; 12 mos.

The research consists of two parts:

1. The Design of Efficient Algorithms. The researcher and his students are studying the best known algorithms for various combinatorial problems such as maximizing network flow, finding maximal (weighted) matching in bipartite (or general) graphs, graph isomorphism in general graphs (or graphs with bounded degree), a membership test for permutation groups and string matching. New efficient parallel algorithms for these problems are sought.
2. Proving Lower Bounds. The researcher and Pavol Duris continue their joint effort in deriving lower bounds for various restricted models of computation.

Cornell University; Vijay Vazirani; *Two Blossoming Paradigms: Algebraic Methods for Computational Combinatoric Problems, and Randomized Reducibilities in Computational Complexity (Computer Research)*; (DCR-8503611); $47,885; 12 mos.

Professor Vazirani is investigating two simple powerful paradigms that benefit the theoretical computer science community:

1. Randomization (in the area of efficient algorithms); and
2. Completeness (in computational complexity).

In particular, he is exploring the use of:

1. Algebraic methods for solving computational combinatorial problems; and
2. Randomized reducibilities for determining the intrinsic computational complexity of natural problems.

A likely outcome of this research is that it will successfully use seldom employed algebraic methods to solve combinatorial problems; for instance, randomized reducibilities will be used to prove that certain "non-number theory" problems are intractable.

Cornell University; Juris Hartmanis; *Computational Complexity and Structure of Feasible Computations (Computer Research)*; (DCR-8301766 A02); $96,266; 12 mos.

During the last five years, research in computational complexity has revealed an unexpectedly rich structure of the feasible computations; it has yielded new research techniques and pinpointed new problems. The structure of the classic complexity classes, such as LOGTAPE, NLOGTAPE, PTIME, NPTIME, PTAPE, etc., is being studied. In the first part of this research the relations between the existence of sets of different densities in these classes and the collapse of the corresponding higher deterministic and nondeterministic complexity classes are being investigated. For exam-

ple, there exist polynomially sparse sets in NP-P if and only if EXPTIME ≠ NEXPTIME. This result has a variety of applications and has suggested new structural properties and relations between the classic complexity classes which are being investigated. The classic separation problem of these complexity classes by means of quasi-homomorphic reductions is also studied. Finally, the computational complexity of finite objects is studied in order to give a feasible and quantitative refinement and extension of the Kolmogorov-Chaitin randomness concepts.

State University of New York - Albany; Victor Y. Pan; *Efficient Computation for Algebraic and Numerical Problems (Computer Research)*; (DCR-8507573); $52,204; 12 mos

Professor Pan is continuing research that addresses problems in the theory of algebraic and numerical computations. More specifically, he is:

1. Designing and analyzing algorithms for algebraic and numerical computation, which include computation with matrices, integers and polynomials, solutions of systems of linear equations and inequalities and solutions of algebraic equations; and
2. Estimating the computational complexity of these algorithms under different models.

State University of New York - Stony Brook; Alan Tucker; *Perfect Graphs, Circular-Arc Graphs, and Associated Algorithms (Mathematical Sciences and Computer Research)*; (DCR-8301934 A01); $5,000; 12 mos; (Joint support with the Algebra and Number Theory Program - Total Grant $31,050).

In his work on perfect graphs, circular-arc graphs and associated algorithms, Professor Tucker has created a unique blend of the mathematical study of combinatorics and its applications to theoretical computer research. The Strong Perfect Graph Conjecture is one of the most important open problems in combinatorial theory today. It concerns a fundamental graph theoretic duality involving cliques (sets of mutually adjacent vertices) and stable or independent sets (set of mutually non-adjacent vertices). The conjecture has been of particular interest to mathematical and computer scientists because of its relation to the development of perfect codes.

University of Oregon - Eugene; Eugene Luks; *The Computational Complexity of Algebraic Problems and Graph Isomorphism (Computer Research and Mathematical Sciences)*; (DCR-8403745 A02); $20,304; 12 mos; (Joint support with the Algebra and Number Theory Program - Total Grant $29,804).

This research focuses on the computational complexity of fundamental algebraic problems, particularly those involving permutation and matrix groups. A major objective is a subexponential algorithm for graph isomorphism using new techniques which are linking such algorithmic problems to recent advances in finite group theory. This would simultaneously resolve many of the remaining complexity questions for permutation groups. Many of the results are extendible to matrix groups and algebras, where similar algorithmic questions are independently motivated. However, it is expected that some of these are provably hard and evidence to support that viewpoint is being compiled.

Carnegie Mellon University; Ravindran Kannan; *Computational Complexity of Numerical Algorithms (Computer Research)*; (DCR-8416190); $53,430; 20 mos; (Joint support with the Computational Mathematics Program - Total Grant $67,430).

Improved algorithms for computing the Smith and Hermite normal forms of an integer matrix along with necessary empirical improvements are studied. Theoretical improvements and implementations of algorithms for solving the same problems for matrices of polynomials are also studied. The "alternation translation" technique is studied in order to obtain results about nondeterminism and determinism using the polynomial time hierarchy. Finally, several computational problems involving multivariate polynomials are attacked.

Carnegie Mellon University; Ravindran Kannan; *Lattice Algorithms and their Application to Combinatorial Optimization*; (ECS-8418392); $13,306; 12 mos; (Joint support with the Systems Theory and Operations Research Program - Total Grant $26,612).

While integer programming problems are known to be hard (NP-hard) in general, there is a great practical need to solve such problems. In the last decade or so, operations researchers have looked at integer programming problems from a computational complexity viewpoint and Lenstra derived a polynomial time algorithm for any fixed number of dimensions. The approach involved solving a well-defined problem on lattices and then applying the result to integer programs. A lattice is the set of all integer combinations of a finite set of linearly independent vectors (the finite set of vectors is called a basis of the lattice). These results indicate that a better understanding of lattices could yield important advances in solving combinatorial optimization problems. This research concentrates on lattices and has two main aims:

1. To devise more efficient algorithms for lattice problems; and
2. To apply these algorithms to solve integer programming and other combinatorial optimization problems.

Pennsylvania State University; Piotr Berman and George Schnitger; *Topics in Computational Complexity (Computer Research)*; (DCR-8407256); $116,949; 12 mos.

Three major areas of computational complexity are studied: the complexity of parallel algorithms, information transfer, and certain problems concerning sequential complexity. In parallel algorithms, algorithm construction, complexity of rounding problems, and restricted models of parallel computation are studied. In the area of information transfer, probabilistic information transfer, transfer with limited memory, and applications to lower bound proofs are investigated. Sequential complexity is also studied.

Brown University; Robert Sedgewick; *Analysis and Implementation of Algorithms (Computer Research)*; (DCR-8308806 A01); $55,833; 12 mos; (Joint support with the Software System Science Program - Total Grant $91,822).

This research concerns the detailed study of the properties of fundamental algorithms. The research centers on drawing specific conclusions on the best way to solve important practical problems, backed up by careful implementations of basic algorithms, full mathematical analysis, and empirical studies to verify the analysis. At best this type of careful scrutiny often leads to the design of new and more efficient algorithms or to the identification of general techniques for use in studying algorithms. The research contributes to the understanding of important methods of problem solution.

Typically, research activity in this field spans a wide range of areas, from topics such as number theory and complex analysis to topics such as programming languages and computer architecture as they relate to a narrower range of algorithms and problem domains, which are studied in detail (always with an eye toward those techniques which are generally applicable). This type of detailed analysis is reserved for the most important problems. Specifically, the research is developing the mathematical analysis necessary to obtain detailed computational complexity results. This is done in the context of important sorting, searching, and computational geometry algorithms. The research also makes heavy use of tools that allow the visualization of algorithms through the use of high resolution color graphics.

Brown University; Jeffrey S. Vitter; *The Design and Analysis of Combinatorial Algorithms (Computer Research)*; (DCR-8403613); $50,706; 12 mos.

This project is concerned with the design and mathematical analysis of combinatorial algorithms, especially those concerned with information storage and retrieval. Topics being considered include: the effects of parallelism on CPU and I/O performance, the power of geometric information in solving graphics problems, the computational complexity of a general software interface for optical disks, space/time tradeoffs in information retrieval, new methods for random sampling, and hashing.

University of Houston - University Park; Ker-I Ko; *A Complexity Theory of Real Computation*; (DCR-8501226); $28,948; 12 mos.

Professor Ko is studying both computational complexity of real functions and the structure of intractable problems. Professor Ko has become a leading researcher in the field of computational complexity of real functions, developing polynomial level complexity of numerical operations (such as roots, integration and differentiation), and developing the relationship between complexity-theoretic properties and analytic properties of real functions. He is extending the complexity of real functions to include the study of lower bound recursive analysis (rather than average case analysis as done by Smale, Shub, Blum, et. al., or information-centered analysis as done by Traub, et. al.) and investigating the complexity and structure of complex functions and functionals.

In addition, he is investigating the structure of feasible computation, extending nontrivial results on polynomial complexity cores and lowness properties of structurally defined problems.

University of Washington; Faith E. Fich; *Concrete Complexity in Parallel Models of Computation (Computer Research)*; (DCR-8402676 A01); $22,000; 12 mos.

Upper and lower bounds for specific problems in various models of parallel computation are studied. For concurrent-read, concurrent-write, parallel random access machines, the effect of different write conflict resolution schemes are investigated. When concurrent-read, but not concurrent-write, is allowed, attention is focused on the relationship between adaptive and non-adaptive algorithms. Finally, the trade-off between depth and size for parallel prefix circuits is examined.

University of Washington; Victor L. Klee; *Combinatorics, Complexity, and Convexity (Mathematical Sciences and Computer Research)*; (DMS-8508758); $5,000; 12 mos; (Joint support with the Algebra and Number Theory Program - Total Grant $35,000).

Professor Klee has been a leading researcher in convexity theory, combinatorical optimization, and discrete mathematics for thirty-five years. He continues to work with vigor and creativity at the frontier of these subjects. This research is concerned with three problem areas: Complexity of optimization problems, qualitative analysis of linear systems and the geometry of convex bodies. The first of these areas is concerned

with measuring the computational cost of optimization algorithms. This is a very active area of research because of the impact of high speed computers on problems of resource allocation and because of recent algorithm discoveries. The second area originates to large extent in mathematical ecomonics. The third area is less immediately applicable but is related to problems of tilings - another area where there have been recent significant results in the form of unexpected examples.

Parallel, Distributed, and Data Base Models

Stanford University; Zohar Manna; *Temporal Verification and Synthesis of Concurrent Programs (Computer Research)*; (DCR-8413230); $68,159; 12 mos.

Temporal logic is a useful, powerful formalism for the analysis and specification of concurrent and real-time systems. It has been found highly applicable in such diverse areas as the design and verification of communication protocols, hardware design, synthesis of concurrent programs, and data base analysis.

Research aimed at making the temporal formalism into a practical tool is establishing a uniform temporal-logic methodology for specification, verification, development, and automatic synthesis of concurrent systems, and is constructing an experimental system that provides computerized support for these activities.

Stanford University; Christos H. Papadimitriou; *Algorithms, Complexity, and Database Theory (Computer Research)*; (DCR-8320000 A01); $101,236; 12 mos.

Professor Papadimitriou is studying a number of areas in the fields of the theory of algorithms and computational complexity. These include linear programming problems, linear decision tree models, decision making under uncertainty, graph theory, and data base theory. Two types of results are sought: foundational and problem specific. The former is illustrated by the search for a proper framework to study the linear programming problem. In one sense this problem is "polynomial" yet it can be argued that the polynomial solution is really "exponential". Satisfactory resolution of this question is sought. An example of a problem specific research area is the study of graph algorithms where the complexity of questions such as "Does a digraph have an even cycle?" is investigated.

University of California - San Diego; Janos Komlos; *Study of Probalistic and Graph Algorithms (Computer Research)*; (DCR-8505053); $82,516; 24 mos.

Professor Komlos is investigating the following problems:

1. Graph algorithms, including spanning tree algorithms and shortest path problems; and
2. Probabilistic algorithms (both in the sense of Solovay-Strassen and the study of heuristic algorithms), including bin packing, expanders and their use, scheduling and minimal Steiner trees.

These investigations have applications in cryptography as well as in the general analysis and design of algorithms.

University of Southern California; Seymour Ginsburg; *Mathematical Foundations of Databases (Computer Research)*; (DCR-8318752 A01 & A02); $69,473; 12 mos.

This study is concerned with theoretical aspects of data bases. The relational database model is extended in order to take advantage of naturally occurring semantic constraints. The algebraic properties of these constraints are studied and applications to database efficiency and the physical design of database machines are explored.

University of Colorado - Boulder; Harold N. Gabow; *Design and Analysis of Algorithms for Networks, Matroids, and Related Problems (Computer Research)*; (DCR-8511991); $84,000; 24 mos.

Professor Gabow is studying techniques for efficient algorithms, especially combinatorial optimization problems on networks and matroids. He has participated in three exciting advances:

1. He instigated (later joined by Tarjan) the "onion/find" operation implementation in linear time;
2. He was a "prime mover" in the resurrection of old scaling techniques as a general tool (better than traditional methods) for optimization; and
3. With his graduate student, Stallman, he completed work on matroid problems which led to a matroid parity algorithm that sped up Lovasz's algorithm from $O(n^{14})$ to $O(n^{5})$.

A likely outcome of this research is that many of these innovative optimization techniques will be expanded producing both important theoretical and practical techniques for developing faster, better and more "understandable" algorithms.

Harvard University; Michael O. Rabin; *Randomized Algorithms and Concurrent Computations (Computer Research)*; (DCR-8121431 A04); $100,000; 12 mos.

Research is concentrated on three interrelating themes centering around randomized algorithms on the one hand and concurrent processes on the other hand: Randomized algorithms for hash functions and pattern matching problems; algorithms for concurrency control and synchronization; concurrent operations on very large data-structures. The implementation of these algorithms in VLSI is also studied.

Harvard University; John Reif; *Theoretical Investigations in Parallel Algorithms and Complexity of Motion Planning (Computer Research)*; (DCR -8503251); $60,768; 12 mos.

Professor Reif is investigating parallel algorithms and robotic motion planning problems. In particular, his characterization of polylog efficient problems permits investigation of parallel algorithm problems that run in polylog time, but whose processor bounds are at most a polylog factor more than best known sequential time bounds. Such problems include:

1. Breadth first search;
2. Matching;
3. Flow (with bounded capacities);
4. Triconnectivity; and
5. Planarity testing.

Using the concept of "polylog time", Professor Reif is exploring the efficient derandomization of parallel algorithms. Finally, he is devising polynomial sequential time algorithms for robotic motion planning problems (involving minimal movement or dynamic movement) or proving the intractability of such problems.

Massachusetts Institute of Technology; Nancy Lynch; *High-Level Models For Reliable Distributed Computing (Computer Research)*; (DCR-8302391 A02); $149,798; 12 mos.

The specific goals of the project are:

1. To evaluate the nested transaction model as a choice of programming model for distributed systems;
2. To isolate and understand the many problems involved in the implementation of nested transactions; and
3. To improve on the usual presentation of distributed database concurrency control theory by viewing it as a special case of the theory of nested transactions.

University of Michigan - Ann Arbor; Quentin F. Stout; *Algorithms for Mesh-Related Computers (Computer Research)*; (DCR-8507851); $87,792; 24 mos.

Professor Stout has done pioneering work in parallel algorithms on hierarchical mesh-connected architectures, optimal algorithms and computational geometry. The problem areas to be adressed include:

1. Determining geometric properties of black/white images;
2. Properties of gray level or range data images;
3. Graph theory; and
4. Computational geometry with point or polygon input data.

Considerations are based on mesh computers of various dimensions with augmentations done by either adding one processor that can send a message to all others at the same time or stacking levels of successively smaller meshes with communication links between levels.

It is expected that these models will lead to improvements to the mesh and will mesh with broadcasting and pyramid algorithms/machines by trying to retain most of the simplicity of these algorithms/machines, while making them more efficient.

Washington University; Jonathan S. Turner; *Evaluation of the Probable Performance of Algorithms Using Domain Specific Distributions (Computer Research)*; (DCR-8409435); $23,085; 12 mos.

This research is aimed at developing better methods for evaluating the performance of heuristic algorithms for hard combinatorial problems. Both worst-case and probabilistic analysis (when applied generally) often fail to provide an adequate analysis of the efficiency of algorithms. The aim in this research is to do probabilistic analysis of algorithms with the probability distribution chosen according to domain specific information. The probability distributions are being chosen by combining both a theoretical and an experimental component. The distributions chosen by theoretical means are being validated by observations of real problem instances.

Princeton University; Robert Sedgewick; *Analysis and Implementation of Algorithms (Computer Research)*; (DCR-8514922); 4 mos; (Joint support with the Software Systems Science Program - Total Grant $49,173).

This research concerns the detailed study of the properties of fundamental algorithms. The research centers on drawing specific conclusions on the best way to solve important practical problems, backed up by careful implementations of basic algorithms, full mathematical analysis, and empirical studies to verify the analysis. At best, this type of careful scrutiny often leads to the design of new and more efficient algorithms or to the identification of general techniques for use in studying algorithms. The research contributes to the understanding of important methods of problem solution.

Typically, research activity in this field spans a wide range of areas, from topics such as number theory and complexity analysis to topics such as programming languages and computer architecture as they relate to a narrower range of algorithms and problem domains, which are studied in detail (always with an eye toward

those techniques which are generally applicable). This type of detailed analysis is reserved for the most important problems. Specifically, the research develops the mathematical analysis necessary to obtain detailed computational complexity results. This is done in the context of important sorting, searching, and computational geometry algorithms. The research also makes heavy use of tools that allow the visualization of algorithms through the use of high resolution color graphics.

City University of New York - Brooklyn College, Kenneth McAloon; *Arithmetic Theories and Incompleteness Phenomena (Computer Research)*; (DCR-8304788 A02); $37,578; 12 mos.

The incompleteness phenomenon is an important characteristic of first order arithmetical systems. Given any mechanizable axiom system for arithmetic, this phenomenon consists of the existence of statements for which neither the statements nor their negations have proofs within the axiom system.

The problems considered include incompleteness phenomena in arithmetic and hiearchy questions for classes of computations. Methods being used include model theoretic techniques, proof theoretic techniques, and diagonal methods.

New York University; Richard Cole and Chee-Keng Yap; *Techniques for Geometric Retrieval and Related Problems in Computational Geometry (Computer Research)*; (DCR-8401633 A01); $61,500; 12 mos.

Computational techniques for geometric retrieval problems are investigated. Three techniques are investigated - invariant regions, coherence, and a form of divide and conquer. Problems studied include k-null problems in computational geometry and motion planning problems in robotics.

State University of New York - Albany; Dan E. Willard; *Data Structures for Predecessor, Geometric, Nonprocedural, Semaphore, and Sequential Retrieval Problems (Computer Research)*; (DCR-8412447); $29,982; 12 mos.

Data structures to accommodate fast searching are one of the most important topics in computer science. This research project deals with several open problems in this topic:

1. Matching upper and lower bounds for the complexity of performing predecessor queries on a set of N positive integer keys which are less than M when the memory space is constrained to O(N) and M is polynomially related to N;
2. Finding tradeoffs between memory space and on-line retrieval time;
3. Extending the approach of applying the batch version of orthogonal range query algorithms to high-level nonprocedural languages; and
4. Investigating ethernet-like semaphores and dense sequential files.

Carnegie-Mellon University; Merrick L. Furst; *Resource Trade-Off Models (Computer Research)*; (DCR-8308805 A01); $52,261; 12 mos.

The main research topic addressed concerns lower time bounds on multi-party protocols. Such protocols arise in distributed systems in which a number of parties (usually system processes) must reach a consensus under the constraints that individually none of the parties has all the information necessary to reach the consensus, yet the sum of the information available to all the parties is sufficient to reach the consensus. Such situations arise, for example, in trying to determine if two processes are simultaneously trying to access the same data items within a distributed environment.

Professor Furst's contribution to this problem was to note a connection between the problem and some deep results in combinatorics involving extensions of Ramsey's theorem. He was able to derive some nontrivial lower bounds for the problem (i.e., lower bounds on the amount of information that must be exchanged to reach a consensus). He has identified several open problems and is attempting to solve them.

Brown University; John Savage; *VLSI Algorithms and Analysis (Computer Research)*; (DCR-8306812 A01); $34,881; 12 mos; (Joint support with the Computer Engineering Program - Total Grant $69,762).

The design of semiconductor chips with thousands or million of active elements is a task in which the management of complexity is a central problem that must be addressed, and one that requires attention to methodology and to performance of algorithms.

Professor Savage's research concentrates on two principal topics:

1. Heuristic algorithms for silicon layout; and
2. Theoretical limits on the performance of parallel and VLSI algorithms.

The first topic seeks to improve a methodology for silicon compilation that produces chip layouts from functional descriptions. The second topic seeks to understand limits on performance of parallel and VLSI algorithms when area, space, time, external space, and inequalities are sought.

Memphis State University; Edward Ordman; *Synchronizing Processes with Partial Simultaneity (Computer Research)*; (DCR-8503922); $39,188; 24 mos.

Professor Ordman is applying methods from graph theory to find more general methods for the implementation of synchronization algorithms for a set of com-

puter processes having limited communication. In particular, the problems addressed include:

1. Cliques that contain n processes of which only one can run;
2. Reader/writer problems where many readers can run, but any writer excludes all other processes; and
3. The dining philosophers problem that is modeled using cycles with adjacent processes.

The results and solutions derive from applications of threshold graphs and coverings by threshold graphs. A likely outcome of this research is that many problems (with synchronizing simultaneous use of computer systems) will be solved and refined.

University of Washington; Richard E. Ladner; *Distributed Algorithms and Complexity (Computer Research)*; (DCR-8402565 A01); $94,139; 12 mos.

Problems in the theory of distributed computing are investigated. These problems include the problem of coordinating a multiple access channel and the development of a game theoretic component as part of a mathematical framework for distributed computing theory.

University of Washington; Walter Ruzzo; *Computational Complexity Theory For Highly Parallel Systems*; (ECS-8306622 A01); $25,000; 12 mos; (Joint support with the Computer Engineering Program - Total Grant $50,000).

The technological capability exists to build very highly parallel computer systems. The technology to effectively use such systems is very much less well-developed. The goals of this research project are to develop theoretical foundations for the design and analysis of algorithms for parallel computer systems, and to apply these principles to elucidate appropriate parallel computer architectures and algorithms.

The five specific areas of research include:

1. Models of parallel computers and their interrelationships;
2. Interconnection networks for parallel machines;
3. Lower bounds or other evidence that specific problems cannot be efficiently solved by parallel machines;
4. Parallel algorithms for a variety of natural problems; and
5. Computational issues arising in VLSI systems.

University of Washington; Martin Tompa; *The Combinatorial Structure of Computations and Symbolic Manipulation (Computer Research)*; (DCR-8301212 A02); $15,393; 12 mos; (Joint support with the Computational Mathematics Program - Total Grant $30,785).

The success of theoretical computer science is due largely to the process of abstracting clean, tractable problems from the morass of realistic detail. Yet there is much left to be desired in the understanding of the structure of computations. For instance, it is not known how to exploit the highly parallel systems that will soon be real. Neither is it known how to exploit randomness (in the sense of probabilistic, or Monte Carlo, computation) to derive efficient algorithms. Even in the setting of ordinary sequential computation, where a reasonably good repertoire of algorithms exists, there are few techniques to show that those algorithms are optimal.

Part of the problem is that not enough of the detail is abstracted away. General computational models like Turing machines or random access machines have sufficient detail to defy analysis.

In this project, computations are abstracted further in order to lay bare the underlying combinatorial structures. With all inessential details stripped away, it is often simpler to discover new efficient algorithms, and demonstrate their optimality (at least within this structured framework) than it is with the details included. The areas being studied are synchronous parallel complexity, sequential time and space complexity, and probabilistic complexity.

University of Wisconsin - Madison; Eric Bach; *Computational Number Theory (Computer Research)*; (DCR-8504485); $81,022; 24 mos.

Professor Bach is investigating the complexity of number-theoretic problems. In particular, he is concentrating on problems in primality testing and factoring for which precise computational complexity and difficulty are still unknown.

His investigations focus on three important subareas of computational number theory that have important connections and applications to cryptography and communications:

1. Applications of analytic number theory to design and analysis of algorithms that test for primality, factor integers, and find lower bounds for complexity;
2. Probabilistic aspects of arithmetic algorithms that explain how to eliminate a "pseudo-random" heuristic; and
3. Relative computational complexity relations among number-theoretic techniques that apply to the practical areas of encryption and pseudo-random number generation.

Mills College; Lenore Blum; *RUI: Computational Complexity of Real Functions and the Loss of Precision (Computer Research)*; (DCR-8417757); $37,790; 12 mos.

In the computational complexity of problems whose natural domain of discourse is the reals, one question is to find the "cost" of obtaining solutions to within prescribed accuracy $E = (1/2)^S$ (or precision $|\log_2 E| = S$). The loss of precision intrinsic to solving a problem, independent of method of solution, gives a lower bound on the cost. It also indicates how realistic it is to assume that basic (arithmetic) operations are exact and/or take one step for complexity analyses. Professor Blum is investigating questions related to the average loss of precision for a wide class of problems, extending previous work on evaluating rational functions and solving linear systems. It is anticipated that insight will be gained toward developing formal models of computation for numerical analysis that are realistic and natural when dealing with questions of computational complexity.

Wesleyan University; Susan Landau; *Algebraic Algorithms and Computational Complexity (Computer Research)*; (DCR-8402173 A01); $11,077; 12 mos; (Joint support with the Algebra and Number Theory Program - Total Grant $16,877).

One of the major mathematical results of the 19th century was the development of a method to determine the solvability of polynomials by Galois. The algorithms for this problem suffered from the defect that they required exponential computation time. Professor Landau developed an algorithm that solved this problem and which required only polynomial time.

She is now studying algorithms for determining the order of Galois groups and for determining the generators of solvable Galois groups. An ultimate goal of this research is a polynomial time algorithm for determining the Galois group regardless of solvability.

Boston University; Leonid A. Levin; *Informational Complexity and Computational Efficiency (Computer Research)*; (DCR-8304498 A02); $54,000; 12 mos.

Informational Complexity Theory (ICT) is an area started from the observation by A. N. Kolmogorov and R. J. Solomonoff that there exists an invariant way to define the notion of Complexity K(x) of a word x - the length of the shortest program producing it. It was found to be fundamentally related to the ideas of randomness, a priori probability, information, inductive inference and others. Research in this area, with emphasis on the questions of computational efficiency, is continuing. Examples of exciting recent achievements related to both areas are the results of M. Blum, S. Micali and A. Yao on pseudo-random number generators. Three particular topics are studied. The first - Linear Programming - does not really need ICT for its justification. It is included because many computational difficulties in ICT are related to handling convex functions. The second problem is the reformulation of ICT concepts in a "limited time environment". The third problem involves an application of ICT to VLSI models.

Harvard University; Harry R. Lewis; *Logic, Circuits and Oracles (Computer Research)*; (DCR-8402489 A01 & A02); $92,081; 12 mos.

Three related areas of computational complexity are studied. First, the study of computational problems in logic is pursued, both decision problems relevant to other computational domains and specific algorithmic problems such as the unification computation. Second, the constant-depth circuit model is studied, both as it relates to logical complexity and with an eye towards extending recently obtained negative results on its computational power. Third, recursion theoretic techniques are applied to relativized versions of some of the basic problems of sequential and parallel computational complexity.

Massachusetts Institute of Technology; John V. Guttag; *Industry/University Cooperative Project: Extensions to the Rewrite Rule Laboratory and Development of Related Theory (Computer Research)*; (DCR-8411639); $25,957; 12 mos; (Joint support with the Software Systems Science Program, and the Division of Industrial Science and Technological Innovation - Total Grant $103,831).

This project is focused on three related topics:

1. Enhancement, maintenance and distribution of rewrite rule laboratories - REVE and RRL;
2. Theoretical and implementation work on the method of proof by consistency; and
3. Theoretical and implementation work on E-unification and generalizations for the Knuth-Bendix procedure.

The MIT-GE collaboration in devising a rewrite rule laboratory is nearly ready to distribute a coherent system for use by a small group of other rewrite rule researchers. This distribution, with the accompanying feedback and maintenance, is a milestone in this particular research area. Availability of a common laboratory system facilitates comparison of algorithms and logical procedures and provides a base for further work on proof by consistency.

Smith College; Joan P. Hutchinson; *RUI: Separator Theorems for Graphs (Computer Research)*; (DCR-8411690 A01); $15,352; 12 mos.

Graph theory and theoretical computer science have enjoyed a rich cross-fertilization. Both disciplines study the same structures sometimes for the same reasons, sometimes for different ones. One area of mutual benefit has been in the development of separator theorems for graphs and in the design of related recursive algorithms, using the technique of "divide and conquer". A separator is a small set of vertices of a graph whose removal disconnects the graph and leaves all components small and of roughly equal size. When such a separator exists and can be found efficiently, other properties of the graph (e.g., a coloring or a maximum independent set) can be determined for the smaller components and pieced together in the original graph. This research focuses on families of graphs that have good separators, on variations in the kinds and characteristics of separators and on the applications of the resulting separator theorems. The research adds a new dimension to that currently being pursued in the mathematics and computer science departments of Smith College. Students majoring in either discipline will be selected for summer projects on the algorithmic and theoretical aspects of the research.

Williams College; Kim B. Bruce; *RUI: Models of Typed and Untyped Lambda Calculus (Computer Research)*; (DCR-8402700 A01); $20,747; 12 mos.

Research is continued on the semantics of typed and untyped lambda calculus. The definition of environment models of the second-order lambda calculus and the related soundness and completeness theorems is refined. Combinatory models for this language are investigated. Similar work is done with the typed lambda calculus. Special effort is put into models of the typed lambda calculus where recursive domain equations must be satisfied. Concrete algebraic methods for solving recursive domain equations are also developed. Finally type assignment for the second-order lambda calculus is investigated.

University of Michigan - Ann Arbor; Andreas R. Blass and Yuri Gurevich; *Computer Science and Mathematical Logic (Mathematical Sciences and Computer Research)*; (DMS-8101560 A02); $20,000; (Joint support with the Topology and Foundations Program - Total Grant $40,000).

The focus of this research is in the area of topos theory, an algebraic approach to the foundations of mathematics which differs radically from the usual foundation research undertaken by logicians.

This supplement provides support for Saharon Shelah of the Hebrew University in Jerusalem to visit the University of Michigan and work with Professor Blass in this area.

University of Michigan - Ann Arbor; Kevin J. Compton; *Computational Problems in Finite Model Theory and Combinatorics (Computer Research)*; (DCR-8418524); $23,344; 24 mos.

This research project deals with two types of problems. The first concerns questions in finite model theory which are closely related to computational problems. The major components of this research are consideration of combinatorial properties expressible in logics more appropriate to asymptotic problems than first order and monadic second order logic; consideration of classes of structures which have figured prominently in combinatorics and computer science; and development of Ehrenfeucht game and analytic techniques for these logics and classes. The second concerns algorithms for the computation of fixed points in ordered semirings. This may be seen as a generalization of the well-known path problems in combinatorics. Results here should have applications to probabilistic computations (in particular, for Markov chains), formal language theory, and combinatorics.

University of Michigan; Yuri Gurevich; *Logic, Complexity, and Semantics (Computer Research)*; (DCR-8503275); $111,030; 24 mos.

Professor Gurevich is applying methods from logic to:

1. Develop logic languages tailored for computational complexity that permit programming exactly the functions of a given complexity class;
2. Study decision and optimization problems for random structures; and
3. Develop uniform families of finite machines and dynamical structures with the bounded resources to permit a "natural" treatment of bounded resources in semantics.

Although fairly mathematical, all of these areas are of importance to the theoretical and foundational areas of computer science.

City University of New York - Baruch College; L. A. S. Kirby; *Models of Arithmetic (Computer Research and Mathematical Sciences)*; (DMS-8308412 A01); $11,318; 12 mos; (Joint support with the Topology and Foundations Program - Total Grant $22,636).

Models of axiomatic systems are mathematical structures which satisfy the axioms and rules of inference of those axiomatic systems. For many axiomatic systems there exists more than one mathematical structure that satisfies the axioms. The intended structure (the one that the axioms were supposed to describe) is known as the standard structure or model; the others are known as non-standard models.

Professor Kirby is investigating models of the axioms of Peano arithmetic (that is, axioms intended to describe ordinary arithmetic). The study of these non-

standard models leads to natural statements equivalent to Godel's incompleteness theorem for arithmetic. The study also leads to results concerning the provability of important computational complexity problems, such as the P = NP question.

City University of New York - Brooklyn College; Rohit J. Parikh; *RUI: Workshop in Logics of Programs, Brooklyn, New York, June 17-19, 1985 (Computer Research)*; (DCR-8502839); $12,780; 4 mos.

A workshop on Logics of Programs was held at Brooklyn College of the City University of New York on June 17-19, 1985. The purpose of this workshop was to bring together proponents of the various schools in the logics of programs, program verification, and program semantics. Topics discussed included Hoare logic, Dynamic logic and its offshoots, Temporal logic, Algorithmic logic, Semantics and proof systems for probabilistic and concurrent computation, and Denotational semantics.

City University of New York - Brooklyn College; Rohit J. Parikh; *RUI: Applications of Logic to Computer Science (Computer Research)*; (DCR-8504825); $70,466; 12 mos.

The primary purpose of this project is to continue work in the logic of programs, especially concurrent and probabilistic computation, but at the same time to explore the possibilities of developing in a precise and rigorous way, related logics suited to artificial intelligence and the representation of human reasoning. Work done under this project is expected to be highly technical in nature consisting of mathematical results and the occasional development of experimental software.

Clarkson University; James F. Lynch; *Problems in Finite Model Theory (Computer Research)*; (DCR-8402206 A01); $19,898; 12 mos.

The objective of this project is to obtain further results in two areas of finite model theory. The first area is concerned with the asymptotic probability that a random finite model satisfies a given first-order sentence. The principal investigator is studying the applicability of several new combinatorial methods for proving that sentences in certain first-order languages always have asymptotic probabilities.

The second area is definability of properties on finite models: given a finite model for each natural number n, and a set of relations on n, how large must a first-order sentence be to define the set of relations on the model? Professor Lynch is studying definability on models with a successor relation.

Columbia University; Joseph F. Traub; *Information and Complexity (Computer Research and Information Science)*; (DCR-8214322 A02); $59,747; 12 mos; (Joint support with the Information Science Program - Total Grant $75,747).

Research is continued on the solution of problems which cannot be exactly solved at finite cost, that is problems which can only be solved with uncertainty. The power of the approach stems from analysis at the information level rather than at the algorithm level.

Most problems can only be solved with uncertainty. This is true of problems from fields as diverse as physical science, biological science, economics, engineering, statistics, statistical decision theory and mathematics. Even the problems which could be solved exactly are sometimes solved probabilistically or approximately, that is, with uncertainty, to decrease their complexity.

Among the new topics studied are: Average Case Analysis, Parallel Algorithms, Value of Information, Distributed Computation, Smoothness and Complexity, Complexity of Partial Differential Equations, Complexity of Linear Programming,and New Applications.

General Electric Research and Development Center; David R. Musser; *Industry/University Cooperative Project: Extensions to the Rewrite Rule Laboratory and Development of Related Theory (Computer Research)*; (DCR-8408461); $12,571; 12 mos; (Joint support with the Software Systems Science Program, and the Division of Industrial Science and Technological Innovation - Total Grant $50,284).

This project is focused on three related topics:

1. Enhancement, maintenance and distribution of rewrite rule laboratories - REVE and RRL;
2. Theoretical and implementation work on the method of proof by consistency; and
3. Theoretical and implementation work on E-unification and generalizations for the Knuth-Bendix procedure.

General Electric Research and Development Center; David R. Musser; *International Conference on Rewriting Techniques and Applications, Dijon, Burgundy, France, May 20-22, 1985 (Computer Research)*; (DCR-8506157); $6,340; 6 mos.

The First International Conference on Rewriting Techniques and Applications (RTA-85) focused on the growing interest in the theory and applications of term rewriting techniques. These applications include automated theorem proving, computer algebra, unification and matching algorithms, functional and logic programming, data type implementation and validation, program specification and verification, program transformation, and synthesis of programs. A major stimulus of interest in the use of rewriting techniques in these applications has been the recent progress in de-

veloping completion procedures for term rewriting systems (which extend a given system into one which produces canonical forms of terms and which therefore can be used to decide equations).

University of Pittsburgh; Robert P. Daley; *Computational Complexity and Inductive Inference (Computer Research)*; (DCR-8402002 A01); $36,408; 12 mos.

Three topics are investigated:

1. An axiomatic approach to the computational complexity of inductive inference;
2. The complexity of concrete procedures for inductive inference, and trade-offs between the complexity of such procedures and various criteria for successful inference; and
3. The differences in complexity between non-deterministic and deterministic computations which use little or no workspace.

University of Wisconsin - Madison; Deborah A. Joseph; *Computational Complexity and Geometric Algorithms (Computer Research)*; (DCR-8402375 A01); $41,956; 12 mos.

Several problems relating to the foundations of computer science and the design and analysis of algorithms are studied. These problems fall into two research areas, the goals of which are to provide additional insight into or answer the following very general questions:

1. What are reasonable formal systems and proof techniques for resolving separation questions concerning complexity classes?
2. In what situations can efficient algorithms be found for motion planning and collision avoidance problems that arise in robotics?

Automata and Language-based Models

University of California - Santa Barbara; Ronald V. Book; *Applications of Combinatorics on Words (Computer Research)*; (DCR-8314977 A01 & A02); $48,793; 12 mos.

Term rewriting systems have been used in a number of areas of computer science, particularly in automated theorem proving and programming language research. Such systems transform strings of characters to strings of characters by rewriting part or all of the original string according to simple rules.

Professor Book is investigating several applications of term rewriting systems. These include studying restricted systems with computationally easy properties, studying applications to context free languages, studying applications to tree-manipulation systems, and investigating network communication protocol applications.

University of California - Santa Barbara; Ronald V. Book; *Computational Complexity Theory (Computer Research)*; (DCR-8312472 A01); $45,948; 12 mos.

The major open question in theoretical computer science is the P = NP question. The P = NP question asks whether problems solvable by nondeterministic programs whose running times increase as a polynomial of the length of their inputs (NP) are also solvable by deterministic programs with polynomial running times (P). The effect of nondeterminism is to allow alternatives in a problem to be explored simultaneously. While no proof has yet been found, most researchers believe that P ≠ NP.

The reason the P = NP question is important is twofold:

1. The problems in P are considered to be "computationally feasible".
2. Representative problems in NP (known as NP-complete problems) constitute an important class of optimization problems.

All direct approaches to proving P ≠ NP have failed. These approaches have concentrated on using "diagonalization" arguments developed by Cantor in the 19th century in an attempt to construct a problem in NP that is different from all problems in P. One of the reasons for the failure of this method is the fact that if a diagonalization proof could show that P ≠ NP, then the same proof could be used to show that P ≠ NP in the presence of "oracles" where an oracle is an extension of the computational power of the programs used to solve the P and NP problems. In fact there are oracles such that both P = NP and P ≠ NP are provable.

This research concentrates on studying the P = NP and related questions in the presence of oracles. The ultimate goal of the research is to gain an understanding of the computational power of nondeterminism.

University of Colorado - Boulder; Andrzej Ehrenfeucht; *New Directions in Language Theory (Computer Research)*; (DCR-8305245 A02); $38,100; 12 mos.

Research is conducted in several areas of formal language theory. The research topics studied include combinatorial properties of formal languages, the Post correspondence problem, the theory of codes, selective substitution grammars, and graph grammars. The main object of the research is to gain further insight into the basic structure of formal languages - addressing both generation questions and decomposition ques-

tions. The research on graph grammars may also have important applications in other areas of computer science.

Iowa State University; Alan L. Selman; *Complexity of Feasible Computations (Computer Research)*; (DCR-8402033 A01 & A02); $52,770; 12 mos.

The computational complexity of feasible computations is studied. The objectives are to reveal the rich structure that individual complexity classes, such as NP, appear to have, and to increase understanding of the structural relations between low-level complexity classes. To these ends several techniques are applied.

Properties of polynomial time-bounded reducibilities on NP, and consequences of P = NP and similar hypotheses, are investigated and compared. Relativization techniques are applied, and fundamental issues are raised concerning the import of results about classes relativized to oracles.

Applications of these studies are being made to problems in cryptocomplexity. Appropriate complexity measures are being identified according to which public-key cryptosystems are to be hard, and then investigations are being made to determine whether hard public-key cryptosystems can exist.

State University of New York - Albany; Harry B. Hunt; *The Complexity of Multi-Set Descriptors, Difference Equations, and Algebraic and Recursively Presented Structures (Computer Research)*; (DCR-8403014 A01); $44,369; 12 mos.

Professor Hunt is investigating the following problems:

1. Multi-set descriptors, unambiguous language descriptors, and difference equations;
2. The complexity of algebra or structures such as lattices, semi-rings, rings and regular algebras; and
3. The complexity of analyzing recursively presented structures.

Particular problems being considered include the following:

1. Finding minimal representations, testing equivalence of Multi-set descriptors, the algorithmic analysis and solution of difference equations;
2. The development of relational algebras and relational calculi for logics other than two-valued logic, the determination of particular algebraic structures for which good algorithms exist for solving systems of simultaneous linear equations; and
3. The complexity of very elementary and basic problems for various classes of recursively presented structures, especially that of determining extraneous productions or state-transitions.

State University of New York - Albany; Richard Stearns; *The Complexity of Unambiguous Representations (Computer Research)*; (DCR-8303932 A01); $113,318; 24 mos.

Professor Stearns is investigating decision problems concerning unambiguous language representations that can be proven "impractical" (i.e., either complex or undecidable). For example, testing equivalence of regular expression is PSPACE complete and yet the same test for unambiguous regular expressions is polynomial. Building on recent successes with unambiguous representations these techniques can be extended to other problems involving unambiguous representations, multi-set representations, and computational algebra.

Computer Research Equipment

University of Illinois - Urbana; Gerald DeJong; *Equipment for Computer Research*; (DCR-8504823); $37,455; 12 mos.

Several existing workstations are being upgraded in support of research in the following areas:

1. Intelligent robotics,
2. Explanation-based learning,
3. Natural language processing,
4. Cognitive modeling,
5. Computer vision, and
6. Computer architecture.

Indiana University - Bloomington; Edward Robertson; *Equipment for Computer Research*; (DCR-8505635); $99,337; 12 mos.

A color graphics workstation and data acquisition system are being obtained in support of research in the following areas:

1. Speech recognition,
2. Diagrammatic information processing in neural arrays,
3. Pattern recognition over 2D-neural arrays using Fourier techniques, and
4. Multiple communicating cognitive agents maintaining propositional attitudes about each other.

Purdue University; John Rice; *Equipment for Computer Research*; (DCR-8505068); $184,500; 12 mos.

A multiprocessor computer system is being obtained in support of research in the following areas:

1. Language for parallel computation,
2. Partial differential equation (PDE) computation, and
3. Multiprocessor algorithm analysis.

University of Southwestern Louisiana; Terry M. Walker; *Equipment for Computer Research*; (DCR-8507014); $52,275; 12 mos.

A microprogram development system is being obtained in support of research in the following areas:

1. Design and implementation of integrated services digital networks, and
2. The development of an integrated set of tools to describe, simulate, and compile code for a class of horizontally programmed architectures.

Brandeis University; Jacques Cohen; *Equipment for Computer Research*; (DCR-8501668); $74,370; 12 mos.

A Lisp machine and workstation upgrades are being obtained in support of research in the following areas:

1. Microanalysis of programs,
2. Compiler prototyping using logic programming,
3. Text compression,
4. Image compression, and
5. Massively parallel knowledge representation and processing.

Tufts University; George Cybenko; *Equipment for Computer Research*; (DCR-8505634); $67,500; 12 mos.

A novel multiprocessor is being obtained in support of research in the following areas:

1. Parallel algorithms for matrix algebra and signal processing;
2. System software support for multiprocessor architecture; and
3. Automatic verification of concurrent programs.

University of Massachusetts - Amherst; David A. Carlson; *Equipment for Computer Research*; (DCR-8505499); $105,413; 12 mos.

An interconnected set of workstations is being obtained in support of research in the following areas:

1. VLSI design research, and
2. Fault tolerant processor arrays for VLSI.

City University of New York - Brooklyn College; David M. Arnow; *Equipment for Computer Research*; (DCR-8506210); $35,406; 12 mos.

A minicomputer system is being expanded in support of research in the following areas:

1. Applied logic and complexity theory,
2. Multidimensional dynamical system analysis of the vestibular ocular reflex with applications to robotics;
3. Data base management systems;
4. Parallel Monte Carlo methods for scientific computing.

Cornell University; David Gries; *Equipment for Computer Research*; (DCR-8505451); $181,125; 12 mos.

A high-performance computer system is being obtained in support of research in the following areas:

1. Fault-tolerant distributed computing,
2. Database system performance,
3. Code optimization for parallel architectures, and
4. VLSI circuit design and layout.

New York University; Olof Widlund; *Computer Research Equipment*; (DCR-8504821); $139,776; 12 mos.

Several high-resolution color workstations are being obtained in support of research in the following areas:

1. Computational geometry,
2. Design of VLSI-based systems,
3. Computer vision,
4. Dextrous manipulation and eye-hand coordination, and
5. Adaptive methods for compressible flow.

Rensselaer Polytechnic Institute; Joseph E. Flaherty; *Equipment for Computer Research*; (DCR-8504487); $80,925; 12 mos.

A large minicomputer is being obtained in support of research in the following areas:

1. Adaptive numerical methods for partial differential equations,
2. Algorithms for computer algebra,
3. Mathematical models for biological cell motion, and
4. Analysis of 2-D gels for estimating the rate of mutation in human germinal cells.

State University of New York - Stony Brook; Alessandro Zorat, Alessandro Giacalone, Jieh Hsiang, Scott A. Smolka, and Mandayam Srivas; *Equipment for Computer Research*; (DCR-8504838); $121,796; 12 mos.

A collection of microcomputers supporting LISP is being obtained in support of research in the following areas:

1. A term rewriting theorem proving environment,
2. Experiments in formally based design environments with two-dimensional user interfaces,
3. Mechanical aids for the specification and verification of concurrent/distributed systems,
4. A PROLOG environment for developing and reasoning about data types, and
5. Fault tolerant VLSI structures.

Pennsylvania State University; Mary Jane Irwin and Robert Owens; *Equipment for Computer Research*; (DCR-8504847); $100,000; 12 mos.

The addition and upgrading of several high capability workstations are being accomplished in support of research in the following areas:

1. Reduced VLSI architectures,
2. Expanded VLSI architectures, and
3. Reasoning about knowledge and belief.

University of Utah; Gary Lindstrom; *Equipment for Computer Research*; (DCR-8504826); $123,600; 12 mos.

A multi-node multiprocessor computer system is being obtained in support of research in the following areas:

1. Multiprocessor emulation,
2. Distributed simulation strategies,
3. Distributable systems for information handling,
4. Distributed sensing and control, and
5. Multiprocessing LISP systems.

University of Virginia; Robert Cook; *Equipment for Computer Research*; (DCR-8504848); $60,770; 12 mos.

A network of workstations is being obtained in support of research in the following areas:

1. Fast algorithms for multi-net wiring problems,
2. Distributed programming and operating system technology,
3. Rapid development of high-quality language translators,
4. Tolerance of processor failure in distributed systems written in Ada,
5. Distributed simulation as an aid to VLSI design,
6. Rapid software prototyping, and
7. Real-time, fault-tolerant local area networks.

Coordinated Experimental Research

Experimental Computer Research

University of Arizona: David R. Hanson: *A Programming Systems Laboratory (Computer Research)*; (DCR-8320138 A01); $679,548; 12 mos.

This project has aided in the establishment of a major experimental facility for research in programming languages and systems. The goal is to develop a comprehensive computing laboratory to support projects in programming languages, software validation, distributed computing, and programming environments.

Research in programming languages is directed towards accelerating the design and implementation of very high-level languages, including the Icon programming language and new functional programming languages. Validation research is concentrated on numerical software testing and on integrated testing tools based on program mutation techniques. In distributed computing research, efforts focus on operating systems, concurrent languages, fault-tolerant systems, and design and measurement tools and techniques. Finally, research on the environmental aspects of programming systems includes studies on the automatic generation of editing systems and user interfaces and on the farther development of an integrated command and programming language.

University of California - Los Angeles: Gerald Estrin, Leonard Kleinrock, Gerald Popek, Walter Karplus, and Algirdas Avizienis; *An Advanced Network Environment for Distributed Systems Research (Computer Research)*; (DCR-8121696 A06 & A07); $797,489; 12 mos; (Joint support with the Department of Transportation - Total Grant $932,489)

This award has aided in the establishment of a major experimental facility to support a program of research on advanced distributed computing systems and their applications. The facility includes a local-area computer network containing a small number of powerful computers, a large number of flexible minicomputers, and special purpose peripheral equipment.

This experimental facility supports investigations of the applications of distributed systems, as well as studies of the fundamental architectural issues of such systems. Results in both these directions are returned to the facility itself in the form of evolutionary enhancements. The facility is at once both an object of and a tool for advanced experimental computer science research.

A variety of research projects are underway. The research program encompasses: distributed operating systems and distributed databases; performance of new distributed architectures and systems; distributed simulation of network architectures; innovative architectures and fault tolerance; performance evaluation of multiprocessor systems; system design tools and user application interfaces.

University of Colorado - Boulder: Leon Osterweil and Robert Schnabel; *A Facility for Research in Numerical Computation and Software Environments*; (DCR-8420944); $601,830; 12 mos; (Joint support with the Office of Advanced Scientific Computing - Total Grant $676,830).

Over the past decade, researchers in the Computer Science Department at the University of Colorado at Boulder have carried out important research in the areas of numerical computation and software environments. This project is aiding in the establishment of a new experimental computational facility consisting of a loosely coupled network of scientific workstations to enable the expansion of these research areas through use of such new hardware and software systems as high-performance workstations, bit-mapped graphics displays, megahertz communication links, attached array processors, networks, and distributed databases. In addition to allowing significant expansion of work in numerical computation and software environments, the facility serves as an effective system for coordinating the projects of the researchers, opening the prospect of mutually enriching research interactions which might be otherwise impossible.

Yale University: Martin H. Schultz, Alan J. Perlis, and Roger C. Schank; *An Attached Processor Systems Laboratory (Computer Research)*; (DCR-8106181 A04); $465,863; 12 mos.

This award has aided in the establishment of a major experimental facility to support research in numerical analysis, scientific computing, algorithms, array processing, computer architecture, artificial intelligence, natural language processing, and systems programming. These new resources enhance the computational capability and overcome address-space limitations which previously restricted natural language research. They also added a powerful array processor to support scientific computing and specialized processors for the LISP programming language. All these resources

communicate over a high bandwidth network, while terminal access is provided through a switching network.

The local network of sophisticated computing resources supports cooperative research among three major research groups: the systems programming group, the scientific computing group, and the artificial intelligence group. The scientific computing group concentrates on new algorithms and architectures, such as attached processors, to find ways to dramatically increase the scale of problems which can be considered for computer solution. The artificial intelligence group pursues studies of integrated understanding and memory by modeling human memory organization and search strategies. Besides pursuing interests in distributed computing, the systems programming group makes contributions to the others by bringing their experience to bear on the problems of supporting attached processors and implementing interactive problem understanding environments.

University of Illinois - Urbana; Charles W. Gear, C. L. Liu, and William J. Kubitz; *A Computer and Software Systems Design Laboratory*; (DCR-8105896 A05); $719,985; 12 mos.

This award aided in the establishment of a major experimental facility to support research and development of computer aids to the design of computer and software systems. The project is resulting in the development of methodologies, techniques and software for supporting various aspects of computer science research. Joint studies by computer science researchers of various backgrounds are being facilitated.

The facility expands on previous departmental resources adding three major computational resources and supporting a high bandwidth network to interconnect the previous and new resources. Highly interactive access is provided through user interfaces and a switching network. Terminals can be directly connected to a particular resource through the switching network while user and process communication is maintained over the high speed network.

A variety of research projects are underway in four major areas: theoretical foundations, algorithm implementation, support software, and run-time systems. The topics include: theory of computation, creation of scientific software, VLSI technology, computer architecture, compilers, databases, inductive inference and conceptual data analysis, and software systems research. The enhanced facility serves as a catalyst for cooperative and productive research.

Purdue University; Douglas E. Comer, J. Tim Korb and Walter Tichy; *High-Level Network Protocols (Computer Research)*; (DCR-8219178 A02); $175,358; 12 mos; (Joint support with the Special Projects Program - Total Grant $275,358).

This research centers on three major problems occurring in distributed and computer networks. The first area of research is on algorithms and data structures to support electronic communication. As electronic mail services provide more information flow, ways to manage this information must be found. The second project concerns integrating the personal workstation into a general network where the computational resources of the entire network are available to the user. The final project deals with the problem of managing data concerning project development in a general distributed environment. As automation aids become more available, the activities of many users in a network which has both shared and independent computational and storage resources must be managed and coordinated.

University of Maryland; Gilbert W. Stewart, Ashok K. Agrawala, Victor R. Basili, Jack Minker, and Azriel Rosenfeld; *A Testbed for Parallel Algorithm Development (Computer Research)*; (DCR-8219507 A02); $825,000; 12 mos.

This project has aided in the establishment of a major experimental facility for the development and testing of parallel algorithms. The facility has been developed around the highly-parallel, ring-structured, ZMOB multicomputer. The project includes the development of necessary hardware interfaces, hosts, and software systems to support the ZMOB machine.

The principal areas of research on parallel algorithms include computer vision, problem solving, and numerical analysis. Some research in distributed systems and language is also being conducted in addition to the creation of a software environment in a number of research areas, particularly in artificial intelligence and numerical analysis, and involves collaborative efforts among a number of research groups within the department.

University of Massachusetts - Amherst; Victor R. Lesser, Krithivasan Ramamritham, and Edward M. Riseman; *A Research Facility for Cooperative Distributed Computing (Computer Research)*; (DCR-8500332); $637,941; 12 mos; (Joint support with the Intelligent Systems Program - Total Grant $1,237,941).

This project is providing an experimental facility for research in cooperative distributed computing based on a tightly-coupled multiprocessor containing between 64 and 128 processing elements. A key aspect of the project is to address issues that arise out of task and problem solving decomposition in which processes cooperate to solve a single integrated problem rather than a set of independent tasks. This project coordinates and integrates large, successful research groups in the areas of intelligent systems (computer vision, robotics and distributed problem solving, adaptive learning networks, neural modeling) and distributed and parallel processing (language primitives, syn-

chronization, concurrency control, network communications, real-time scheduling, load balancing, debugging, protection and reliability).

University of Minnesota; Sartaj Sahni; *Experimental Research in Computer Algorithms*; (DCR-8420935); $205,946; 12 mos; (Joint support with the Office of Advanced Scientific Computation - Total Grant $411,891).

This project is concerned primarily with experimental research utilizing supercomputers in the design and analysis of algorithms and applications software to serve projects in combinatorial algorithms, optimization, systolic systems, and linear algebra. Expertise in each of these research areas currently exists at the University of Minnesota. The research emphasizes the design of algorithms for problems for which the currently best known algorithms are unable to solve problems of desired sizes using a reasonable amount of computer resources.

A major portion of the work is being carried out on facilities of the Supercomputer Institute at the University of Minnesota. This institute currently operates a Cray 1, a single processor Cray 2 installed in June 1985, and a four processor Cray 2 to be installed in December 1985. In addition, some work is being carried out on certain tightly coupled parallel computers including the HEP and MPP via the University's ARPA net connection and some program development and related activities are being carried out on workstations acquired under this project.

Duke University; Donald J. Rose; *VLSI Computing Structures and Design Methods and Interactive Computer Graphics (Computer Research)*; (DCR-8309911 A03); $328,380; 12 mos.

This award has aided in the establishment of a major experimental facility being shared between researchers at Duke University and the University of North Carolina at Chapel Hill. The facility supports a highly integrated program of research centering on VLSI computing structures and design methods and on interactive computer graphics. The project is also closely tied to the state-funded Microelectronics Center of North Carolina.

The facility includes a variety of computing resources including workstations, graphics equipment, and general purpose computers. The Microsystems Laboratory at the University of North Carolina, used heavily by the Departments of Computer Science at both Duke University and the University of North Carolina for design and testing of VLSI components and systems, has been expanded.

The research is being conducted at three levels in a hierarchy beginning with fundamental VLSI design support and tools, through specialized and general purpose VLSI-based architectures, to the integration of new architectures with software in an interactive computer graphics environment. Each of these areas is closely related. Research in graphics, software engineering, man-machine interfaces, and simulation is providing new ideas for the development of VLSI tools and computer-aided design (CAD) systems. The VLSI CAD systems support research on new highly-parallel architectures and new architectures for computer graphics.

University of North Carolina - Chapel Hill; Donald F. Stanat; *VLSI Computing Structures and Design Methods and Interactive Computer Graphics (Computer Research)*; (DCR-8219306 A02); $878,688; 12 mos.

This award has aided in the establishment of a major experimental facility being shared between researchers at Duke University and the University of North Carolina at Chapel Hill. The facility supports a highly integrated program of research centering on VLSI computing structures and design methods and on interactive computer graphics. The project is also closely tied to the state-funded Microelectronics Center of North Carolina.

The facility includes a variety of computing resources including workstations, graphics equipment, and general purpose computers. The Microsystems Laboratory at the University of North Carolina, used heavily by the Departments of Computer Science at both Duke University and the University of North Carolina-Chapel Hill for design and testing of VLSI components and systems, has been expanded.

The research is being conducted at three levels in a hierarchy beginning with fundamental VLSI design support and tools, through specialized and general purpose VLSI-based architectures, to the integration of new architectures with software in an interactive computer graphics environment. Each of these areas is closely related. Research in graphics, software engineering, man-machine interfaces, and simulation is providing new ideas for the development of VLSI tools and computer-aided design (CAD) systems. The VLSI CAD systems support research on new highly-parallel architectures and new architectures for computer graphics.

Princeton University; Richard Lipton; *Industry/University Cooperative Project: The Massive Memory Machine Project*; (DCR-8420948); $191,816; 12 mos; (Joint support with the Industry University Cooperative Program - Total Grant $255,755).

This project is directed at the investigation of supercomputers with massive amounts of physical memory (tens of billions of bytes). The hypothesis of the research is that such machines may change, in fundamental ways, how certain classes of problems are solved and may lead to orders of magnitude performance improvements. The project is jointly supported

with the Office of Naval Research (ONR) and the Defense Advanced Research Projects Agency (DARPA). Collaboration with Digital Equipment Corporation in the hardware design phases is planned and the project will also be partially supported by the NSF Industry/University Cooperative Research Projects program.

The fundamental goal of the NSF sponsored portion of the project is to validate experimentally the massive memory hypothesis, to explore ways to build such a machine in the future and to investigate selected algorithmic processes for such designs. The validation portion of the project is being carried out on a Large Memory Computer (LMC), a conventional computer with an unconventional but economically feasible amount of memory (either 128 or 256 MB). The second part of the project involves the study of a number of novel architectures for machines with billions of bytes of memory. The architectures are being compared among themselves, as well as with architectures for highly parallel machines, identifying the strengths and weaknesses of each and selecting, from among these choices, a possible optimum choice for such a massive memory machine. The DARPA contract is for the construction by the end of the fifth year of a full-scale Massive Memory Machine prototype capable of supporting billions of bytes of memory. The NSF grant has considerable overlap during the first year but diverges considerably in later years.

Cornell University; Robert L. Constable, Richard W. Conway, David J. Gries, and Alan J. Demers; *A Laboratory for Experiments on the Programming Process*; (DCR-8105763 A04); $573,955; 12 mos.

This award has aided in the establishment of a major experimental facility to support research on the programming process. The investigators are considering a variety of topics including programming environments, the discipline of programming, and language design and implementation. While having distinct goals and clearly separate approaches the high degree of interaction among the researchers is expected to provide results which strongly impact our understanding of programming.

The facility includes a variety of current and new computational resources interconnected by a local area network. While each researcher may access specialized resources for a particular project, all resources are available to each user of the network. This environment provides excellent communication between research projects while matching the appropriate resource to an individual project. Cornell has been characterized as being engaged in highly individualized computer science research in an atmosphere of close rapport. This facility is enhancing the research opportunities while maintaining this highly productive environment.

New York University; Olof Widlund and Jacob T.Schwartz; *Industry/University Cooperative Research: A Laboratory for Advanced Computational and Software Techniques (Computer Research)*; (DCR-8320085 A01); $797,413; 12 mos; (Joint support with the Industry University Cooperative Research Program - Total Grant $1,047,413).

This project has aided in the establishment of a major experimental facility for research, centering on robotics as a focus, uniting key and demanding problems in computer science and engineering and in applied mathematics. The robotics research draws upon and ties together three other areas of current research: software technology, super-speed parallel computation, and scientific computation/numerical analysis.

Research in robotics and computer vision is evolving from initial theoretical studies on collision-free paths, motion planning, hybrid position/force robot control and various problems in which geometry and force both play roles. A basic robotics and vision laboratory has been developed and cooperation with other robotics activities at Cornell and IBM has been established. Research enabled by this project concentrates on servo control and planning algorithms, techniques for processing robot data, and design and development of languages for robotics.

State University of New York - Stony Brook; Jack Heller; *A Data Oriented Network System (Computer Research)*; (DCR-8319966 A01); $1,166,076; 12 mos.

This project has aided in the establishment of a major experimental facility for research on operating systems, programming environments, programming languages, and databases. The design and construction of a data-oriented network system within a network environment as a research facility is a large part of the project. The data-oriented network system exploits the conceptual simplicity of the relational data model to benefit end-users and developers of applications systems by providing a high-level, uniform view of data. The system's facilities include a relational operating system environment, a relational editor, support for logic programming, and a software development environment. Applications of this facility to graphics, VLSI, office automation, and natural language processing are planned. The network software provides location-transparent access to distributed relations and operating system support for distributed programs. The system is a testbed for research into such issues as resource allocation, concurrency control, and multicast communication.

University of Rochester; Jerome A. Feldman; *A Testbed for the Study of Parallel Computation (Computer Research)*; (DCR-8320136 A01); $935,000; 12 mos.

This project has aided in the establishment of a major experimental facility for research on parallel computation centered on a 128-processor multicomputer. The goal is to develop a comprehensive computing laboratory to support projects on massively parallel models, computer vision, knowledge representation, natural language understanding, computer architectures, and systems for parallel computation.

A tightly coupled, parallel machine has been purchased and integrated into the Rochester research network. With a large machine having adequate reliability and software support, it is possible to experiment with truly parallel algorithms. Current research at Rochester is a particularly good base for pursuing this line of inquiry. The researchers have been studying computational models which involve millions of active processing elements. Although these models are not yet practical, they can be effectively simulated on machines with highly parallel structures. The motivation for this research comes from a continuing interest in artificial intelligence, particularly in computer vision. Currently, research in natural language and knowledge representation does not appear so naturally parallel as vision research but preliminary results with distributed systems show promise that algorithms for tightly coupled processors can be found. This environment is an excellent testbed for systems and architecture research and the knowledge gained in a truly parallel environment will impact on these research areas.

University of Pennsylvania; Aravind K. Joshi, Ruzena Bajcsy, and Peter Buneman; *Modeling Interactive Processes: Flexible Communication with Knowledge Bases, Computer Interaction in Three Dimensions (Computer Research)*; (DCR-8219196 A02); $643,117; 12 mos.

This award has aided in the establishment of a major experimental facility to support two main areas of research: flexible communication with knowledge bases and computer interaction in three dimensions. The two areas share several technical aspects and are united by the general goal of modeling an interactive process, whether it is a person interacting with a machine, or a machine with its environment. The latter is relevant to robotics.

Research in these topics involves many established areas in computer science: natural language interfaces to knowledge bases; database systems and methods for interacting with heterogeneous, distributed databases; graphic interfaces for databases; three-dimensional object recognition; representation of visual and other sensor information, in particular, tactile information; integration of sensory information from different modalities; architecture of special purpose machines; movement description and simulation; and a number of related areas.

The facility includes a variety of computing resources including general and special purpose devices. This equipment and the associated personnel support, together with the supporting environment currently existing, enables the investigators to carry out these long-range research programs of technological and scientific significance, and also provides a superb environment for training in experimental computer science.

Brown University; Thomas W. Doeppner; *Third Annual NSF/CER Conference, March 7-8, 1985 Providence, RI*; (DCR-8509741); $31,781; 3 mos.

The Third Annual NSF Coordinated Experimental Research (CER) Conference brought together two or more representatives from each of the then existing eighteen CER grantee institutions along with guests from academia, industry and government. The purpose of the meeting was to have the various schools discuss their research and common problems and to develop joint approaches to solving these problems. One feature of the year's meeting was presentations on their research by the schools most recently awarded CER grants. A second focus was on organized panel discussions of important problems of common interest facing academic experimental computer research programs.

Brown University; Thomas W. Doeppner, John E. Savage, Peter Wegner, and Andries van Dam; *An Integrated Environment for Research in Computer Science*; (DCR-8121806 A03 & A04); $500,173; 12 mos.

This award has aided in the establishment of a major experimental facility to support a program of research in computer science. The facility includes a local-area computer network containing a powerful minicomputer, a large number of personal computers, and special purpose peripheral equipment.

The investigators are developing a computing environment for the effective use of this network of powerful personal computers. At the core of this environment is a powerful set of graphics tools. These tools allow each researcher to customize a graphics-oriented user interface to a specific research project. A database management framework and system underlies and supports the graphics tools.

Research in graphics, databases, and operating systems is key to developing a stable and useful system. The investigators provide the graphics tools to the larger research community for assistance in guiding refinements and revisions. Other research by the investigators in algorithms, artificial intelligence, compilers, document preparation, learning environments and other areas has made good use of existing tools. The interaction between this research and that connected with the facility itself is also leading to improvements and new research directions.

Rice University; Kenneth W. Kennedy, Robert S. Cartwright, John E. Dennis, and J. Robert Jump; *An Experimental Computer Network to Support Numerical Computation (Computer Research)*; (DCR-8121884 A03); $620,580; 12 mos.

This award has aided in the establishment of a major experimental facility to support a program of research on advanced distributed computing systems and numerical computing. The Rice Numerical Network, or R^n, consists of approximately 24 single-user numerical machines equipped with high-resolution bit-mapped screens, a 32-bit central processor, and vector floating-point hardware. It also includes several specialized server nodes and various peripheral devices including a gateway to the CSNET communications network linking the nation's major computer science research centers.

The new facility supports a coherent research program in software systems, computer architecture, and quality numerical software, directed at creating a modern interactive environment for numerical computation.

University of Texas - Austin; Alfred G. Dale, James C. Browne, Donald I. Good, and Abraham Silberschatz; *An Experimental Computing Facility For the Design and Analysis of Reliable High-Performance Distributed Systems (Computer Research)*; (DCR-8122039 A03); $662,300; 12 mos.

This award has aided in the establishment of a major experimental facility that incorporates mid-scale computers with large virtual addressing capabilities and powerful individual workstations linked over a high-speed local network. The facility, together with a professional support staff, provides an environment for the development and experimental use of software for mechanical program and theorem proving, performance measurement and distributed computing. This integrated system design and analysis facility (SDAF):

1. Permits the consolidation of methods for mechanical program verification and performance analysis into an integrated approach to the analysis of distributed systems,
2. Supports extensive experimental validation of formal analysis techniques, and
3. Supports novel applications of formal analysis methods to intelligent user-system interfaces and VLSI design.

University of Utah; Richard F. Riesenfeld and Lee A. Hollaar; *A Laboratory for Computer-Aided Design (Computer Research)*; (DCR-8121750 A04); $446,236; 12 mos.

This award has aided in the establishment of a major experimental facility to support a program of research on Computer-Aided Design (CAD) Systems. The facility includes powerful design workstations based on personal graphics-oriented machines. Included also is a necessary expansion of general computing facilities to accommodate the expanded research effort. Several specialized 3-D output devices provide a basis for testing the CAD techniques and tools developed as part of the research.

There are two main thrusts to the current efforts: Computer-aided 3-D shape design and Very Large Scale Integration (VLSI) design. The ALPHA-1 development project, at the heart of the first research thrust, is capable of supporting a wide range of CAD activities. Moving ALPHA-1 to a distributed environment requires research in operating systems and in workstation support and design. The investigators are also developing true 3-D display of surfaces and Numerical-Control Milling Algorithms. Research in portability and software tools will lead to an enhanced software development environment; in database systems it will lead to an efficient uniform means for accessing, modifying, and sharing design data; and in multiprocessor and network operations it will allow the various tasks in the systems to be partitioned and run on processors best suited to their particular requirements. In integrated circuit design, the main emphasis is on simulation systems improvement, on mapping techniques from high-level representations directly to design, on graphical representations of VLSI designs, and on the testing and diagnosis of integrated circuits.

University of Washington; Edward D. Lazowska; *Integration in Homogeneous and Heterogeneous Distributed Computer Systems (Computer Research)*; (DCR-8420945); $920,868; 24 mos.

This project is centered on the study of distributed computer systems. It has two aspects. The first aspect of the study is to conclude ongoing research concerned with tight integration in homogeneous distributed computer systems — the Eden project. The problems here are well understood although there is considerable disagreement regarding various proposed solutions and further research is appropriate. In this first case, the work focuses on the design of operating system software that allows tight transparent sharing of resources among a collection of essentially identical workstations. The second aspect of the project is to conduct a preliminary study of the problems of accommodating heterogeneity — that is, achieving looser but still significant integration among a collection of heterogeneous computer systems. Here the problem is much less well understood and considerable preliminary work is necessary. Heterogeneity is an inevitable consequence of experimental computer science activity and developing the capability to couple dissimilar machines or subsystems into an existing research environment sharing services such as file access, printing, mail, remote computation, etc. is now being recognized as a key challenge facing the computer science research community. The objectives of

this aspect of the project are to stimulate the sharing of relevant information and specifications, to delineate existing and possible solutions, to differentiate areas requiring development from areas requiring research and to design and build simple prototype solutions to one or more subproblems.

University of Wisconsin - Madison; Lawrence H. Landweber; *A Software-Partitionable Multicomputer: A Testbed For Distributed Processing (Computer Research)*; (DCR-8105904 A06); $957,557; 12 mos.

This award has aided in the establishment of a major experimental facility to support research in databases, distributed processing, operating systems, multicomputer architecture, and numerical analysis. The investigators have implemented and studied a computer architecture that will exploit recent advances in computer and communication technology. The facility includes a multicomputer, a small number of high-performance personal computers, and a support center. The multicomputer implements the target research environment. The personal computers serve as prototypes for program development research projects. The support center computers provide program preparation, simulation, file, and device services.

The multicomputer is a research tool and a research object for a variety of projects. The implementation of the architecture and kernelized operating system is itself a major effort. The prototype system allows the multicomputer to be logically partitioned for independent simultaneous use by different researchers. A large number of projects are supported to examine research issues in the respective areas.

Computer Science Research Network (CSNET)

University Corporation for Atmospheric Research; Leonard C. Romney; *Computer Science Research Network Host*; (DCR-8304539 A02); $228,701; 12 mos.

The Computer Science research Network, CSNET, has been under development by contractors to the NSF for several years. CSNET is a logical network comprised of several physical computer communications networks including ARPANET, Phonenet and X.25-based commercial common carriers such as GTE-Telenet. The objective of CSNET is to make available uniform network services to the entire computer research community with network costs tied to performance and level of service. Among the services offered are electronic mail, file transfer, and virtual terminal access (computer-controlled remote access). The project has concentrated on the development of software implementations of network protocols and on the expansion of Phonenet capabilities. Phonenet is a telephone-based mail relay system which provides electronic mail and transfer services. The CSNET Project is managed by the University Corporation for Atmospheric Research (UCAR) who houses and supports the CSNET Executive Director and the CSNET Executive Committee. Technical operations and the Communications and Information Center (CSNET-CIC) are managed by Bolt, Beranek, and Newman under a subcontract. Substantial support for this project is provided through user fees and dues.

BLANK PAGE

Presidential Young Investigator Awards

Stanford University; David R. Cheriton; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352048 A01); $75,000; 12 mos.

Communication is the key component of distributed and parallel computer systems, being the basic facility that allows the multiple components of the systems to cooperate in sharing information and hardware resources.

This project investigates:

1. Use of communication in highly parallel and distributed computations with particular interest in the benefits of "group" communication paradigms, such as those supported at the hardware level of broadcast networks.
2. The design of communication primitives for parallel computations. Starting from the V kernel IPC design, it is hypothesized that some modest but crucial changes will be required to meet the needs of parallel program structures.
3. The efficient implementation of interprocess communication both on a multiprocessor machine and across a local network joining multiple processors. The feasibility of a special-purpose IPC chip for this purpose as well as the design of intelligent network interfaces for local network communication are being investigated.

Stanford University; John L. Hennessy; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351269 A01); $42,500; 12 mos.

It is anticipated that new computer architectures will use instruction sets that decompose further than existing architectures. Such computers will require improved optimization techniques in compilation so as to produce efficient programs. This project is investigating techniques that can shift computations from run-time to compile-time and from inner loops to outer loops.

These same optimization techniques can be applied to the efficient compilation of applicative programming languages. The efficiency of such languages can be improved by a two part approach: program transformations to eliminate redundant computations and application of optimization techniques to programs during compilation.

Stanford University; Ernst W. Mayr; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351757 A01); $66,500; 12 mos.

The following topics are studied:

1. Paradigms for the design of good parallel algorithms for a variety of parallel architectures.
2. The inherent parallel complexity of algorithmic problems.
3. Tradeoffs between parallel time and communication complexity.
4. Automatic or semi-automatic program transformations for computer-aided design of efficient parallel architectures, and
5. Heuristics for general parallel search problems.

Stanford University; Brian K. Reid; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352015 A01); $92,486; 12 mos.

This project is directed toward the use of programming language techniques to solve problems in manufacturing automation. The long term objective is a language in which engineers can specify manufacturing processes. Such specifications should be archivable and portable between properly automated factories. The available theoretical background for the intellectual support of such a language is limited. Specific research areas are nondeterministic semantics and mechanisms for representing non-hierarchical structures.

University of California-Berkeley; Paul N. Hilfinger; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451213); $25,000; 12 mos.

This project is an investigation of the use of a high-level notation for digital signal processing and of its translation into programs for specialized hardware. The source language, called Silage, is a simple applicative notation that incorporates primitive operators and functions for such basic signal processing operations as delay, interpolation, and decimation. The use of Silage for describing algorithms from a variety of applications is being investigated. A retargetable translator to investigate the issues involved in translating a language at this level both for a (conventional) commercial signal processing chip, and for a customized family of processor chips generated from microcode-like descriptions is also being built.

University of California - Irvine; Richard Taylor; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451421); $25,000; 12 mos.

Assuring that real-time or concurrent software performs its intended function is a difficult task, arguably more so than for serial software. Recent research has developed some promising new techniques for aiding in the checkout of this type of code, including static concurrency analysis, program instrumentation with run-time monitoring, and multi-tasking debugging techniques. Not surprisingly, however, no single technique is adequate for all the demands. What is necessary is the cooperative, mutually supportive application of them all. The primary goal of this project is the further development of these and other techniques, emphasizing their integrated application.

An important recognition guiding the research is that effective investigation and application of integrated techniques requires that the tools embodying the techniques be housed in a programming environment. In so doing duplication of tasks can be avoided, a uniform user interface can be provided, and efficiency can be promoted. The analysis, testing, and debugging tools are thus being developed in concert with a programming environment.

University of California - Los Angeles; Eliezer Gafni; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451396); $55,000; 12 mos.

Three main efforts are being pursued:

1. Design and Analysis of Distributed Algorithms implemented on a Network of Processors. A thorough understanding of basic algorithms of this type is to be achieved before there is a chance for a theory to emerge. This effort concentrates on the time-message tradeoff issue.
2. Distributed Simulated Annealing. The problem of control of a distributed system is inherently distributed. Among the control problems are routing of messages and load management. A distributed version of Simulated Annealing (at fixed temperature) is being employed in order to achieve on line quasi static control.
3. Efficient Algorithms for Multi-Channel Broadcast (MCB) system. The MCB system is an attractive alternative to a single ultra-broad bandwidth channel. It reduces contention and can achieve cumulative high throughput using relatively slow and inexpensive communication components. This effort concentrates on the problem of scheduling communication among processors, given their communication requirements. Centralized as well as distributed solutions are considered.

Yale University; Paul R. Hudak; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451415); $25,000; 12 mos.

Using the lambda calculus as a generic foundation for computer programs, the goal of this project is to investigate semantic analysis techniques to enhance the execution of programs on parallel architectures. General inferencing theories, including ones for higher-order functions, are being developed, and their theoretical power and computational complexity are being characterized. In addition, specific inferencing theories for program properties important to parallel computation (such as function strictness), and specific inferencing algorithms that are practical for use in real compilers are being developed. Since all deterministic programming languages may be cast into the lambda calculus via denotational semantics, the analysis is expected to be general enough for most conventional languages. The culmination of this work will be the implementation of the algorithms in a compiler first for a functional programming language, and ultimately for a more conventional language.

University of Chicago; Theodore A. Slaman; *Presidential Young Investigator Award (Mathematical Sciences and Computer Research)*; (DMS-8451748); $12,500; 12 mos; (Joint support with the Topology and Foundations Program - Total Grant $25,000)

This research is concerned with the study of definability (especially computability) in a mathematical setting and with degrees of undefinability. The research concentrates on:

1. The global structure of the recursively enumerable degrees,
2. Definability in the set of recursively enumerable degrees,
3. The connection between forcing constructions in set theory and effective enumeration in recursion theory, and
4. The foundations of the priority method in recursion theory.

Purdue University; Mikhail Atallah; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451393); $49,000; 12 mos.

Professor Atallah is investigating the design and analysis of algorithms. In particular, the following areas are being investigated:

1. Parallel algorithms: the design of efficient parallel solutions for problems (from computational geometry and artificial intelligence) for which the known sequential algorithms have slow parallel implementations;
2. Computational geometry: problems in computational geometry in a framework that is very likely

to occur in practice (namely when every one of the geometric objects is moving in a prescribed manner);

3. Placement problems: the design of efficient algorithms for solving various placement problems in VLSI layout design (i.e., problems where there is some choice in terms of where to place objects such that some optimality condition is satisfied); and
4. Equivalence of networks: equivalence between various networks of processors.

Johns Hopkins University; Joseph O'Rourke; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351468 A01); $62,500; 12 mos.

Research is progressing along five lines:

1. Algorithms for approximating (by e.g. circumscribing) polygons and polyhedra.
2. Coordinating the movement of several objects in a cluttered environment.
3. Terrain navigation on an undulating surface.
4. Computational geometry on curved manifolds; and
5. The shape of non-inferior regions in multi-objective programs. A monograph is being written on "Art Gallery Theorems and Algorithms."

University of Maryland - College Park; Dana S. Nau; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351463 A01); $62,500; 12 mos.

This research is concerned with several topics involving reasoning and problem solving methods in the field of artificial intelligence.

In particular it is:

1. Exploring an approach to knowledge based problem solving, which is based on a generalization of the set covering problem. This approach has potential application in the development of expert systems for diagnostic problem solving and other areas;
2. Investigating the applicability of expert computer system techniques to a number of automated manufacturing activities including process selection, stock selection, error classification and diagnosis, and the kinds of spatial reasoning that are required for manufacturing planning activities; and
3. Investigating game tree searching. The topics being investigated are the existence of pathology in game trees, the factors governing why pathology occurs, and various alternatives to minimaxing which may avoid the occurence of pathology.

Massachusetts Institute of Technology; Frank T. Leighton; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352349 A01); $62,500; 12 mos.

Problems involving very large scale integration (VLSI) are studied. Two specific topics are investigated:

1. Techniques for solving design-related problems, and
2. Strategies for efficiently utilizing systems of microprocessors.

Problems investigated include: optimal layouts for the shuffle-exchange graph layout theory, three-dimensional VLSI, wafer-scale integration of systolic arrays, channel routing, methodologies for designing fault-free processor arrays, algorithms and networks for parallel computation, and gate array routing.

Rutgers University - Busch Campus; Tom M. Mitchell; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351523 A01 & A02); $89,300; 12 mos.

This research project is concerned with two areas of artificial intelligence: knowledge-based expert systems, and machine learning. One research thrust focuses on developing a knowledge-based consultant system to aid in the design of digital VLSI circuits. This research involves extending current methods for developing rule-based expert systems, by applying them to design tasks, and by integrating rule-based methods with algorithmic methods for dealing with different aspects of the design task. The second research area, machine learning, has focused on developing the LEX system which learns problem solving heuristics by generating and solving practice problems. This research has led to development of a domain-independent technique (goal-directed learning) for generalizing from examples. Research on machine learning includes an attempt to apply these techniques to the task of learning rules for VLSI design.

Columbia University; John R. Kender; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351852 A01); $98,034; 12 mos.

The principal goal is to create and examine representations and algorithms for the detection and confirmation of visual surfaces by a mobile camera. The foundation of the work is the inversion of the mathematics of perspective projection, as applied to modelable textures. Such "shape from texture" work is being pursued experimentally by mounting a camera on a robot arm, digitizing and processing the images, and transforming the two-dimensional data into inferences about local and global three-dimensional surface properties.

Among the investigations being included are the following:

1. Optimal algorithms for edge detection, based on psychophysical observations (Laplacian of Gaussian operations);
2. Cooperation of various low-level visual operations to arrive at an intelligent consensus of surface properties;
3. Inference mechanisms using both optimal interpolatory techniques and experimental control structures (simulated annealing); and
4. Investigations of the applicability of these algorithms to special purpose multiprocessors.

Cornell University; John Gilbert; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451385); $25,000; 12 mos.

Professor Gilbert is investigating the development of combinatorial algorithms which apply to or are suggested by numerical problems. In particular, these investigations extend prior pioneering work. Professor Gilbert is investigating sparse matrix algorithms which include partial pivoting algorithms, sparse null space bases and least square problems and topological graph algorithms, some of which have applications to sparse matrix problems.

University of Rochester; James F. Allen; *Presidential Young Investigator Award (Computer Research)*; (DCR-8351665 A01); $25,000; 12 mos.

Natural language communication is enabled only by the conversants' rich body of knowledge about the topic being discussed, knowledge about the conventions of language itself, and specific knowledge about the other conversants. This project is investigating some basic issues in knowledge representation that are crucial to understanding how communication occurs. In particular, formal models of a naive person's knowledge about actions, space, time, and how they interact in plans are studied. An important subclass of actions are communicative actions.

The results of the work will be published in the artificial intelligence literature. In addition, as the formal models are developed, they are being incorporated as appropriate in ongoing projects to build systems to model natural language communication in dialogues, and problem solving in a dynamic, changing world.

Carnegie-Mellon University; Merrick L. Furst; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352081 A01); $62,500; 12 mos.

Four areas are studied:

1. Computational experiments to discover applications on improvements of Lenstra, et.al.'s shortest vector algorithms;
2. Upper and lower bound proofs for multi-party communication protocols;
3. Semantics-directed code generation; and
4. Computational aspects of group theory.

University of Pennsylvania; Yuen-Wah Eva Ma; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451408); $50,000; 12 MOS.

This research is focused on the area of reconfigurable computer systems, including (i) unifying the study in reconfigurable computer systems, and (ii) designing specific reconfigurable computer architectures which include dynamic computer networks and REPLICA.

Reconfiguration is an important technique for the design of fault-tolerant computer systems. It can also be used to enhance the performance of a system. To date, the study of reconfiguration for fault-tolerance and for performance improvement is divergent. The development of a unified study of both fault-tolerance and performance enhancement in reconfigurable systems is being explored.

A dynamic computer network can be defined as a computer network in which the hosts are mobile. The topology of a dynamic computer network can vary rapidly during system operation. Dynamic reconfigurations must be incorporated in these systems to cope with the variations in their topology. Research in this area focuses on the study of reconfiguration techniques and reliability modeling.

REPLICA is a reconfigurable, partitionable and highly parallel computer architecture for supporting active multi-sensory perception of 3-D objects.. The focus in this area of research includes the design of interconnection networks, memory management techniques, and parallel algorithms.

Brown University; Jeffrey Vitter; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451390); $55,000; 12 mos.

Professor Vitter is investigating a theoretical basis for recovery in programming environments which incorporate high-level recovery packages for user interfaces. Topics being investigated include:

1. Storage and retrieval specifically geared toward obtaining simplified analyses of linear probing and a consolidated treatment of coalesced hashing in order to put things in perspective;
2. Hashing, with lazy deletion;
3. In the area of parallel processing, how to parallelize combinatorial graph problems aand tradeoffs between the amount of parallelism and time; and
4. In the area of basic data structures, deriving theoretical bounds on simple priority queues that seem to perform extremely well in practice and developing better approximation and exact algorithms for problems with geometric constraints (e.g., the trave'ing salesman minimum

spanning tree, and shortest path problems in Euclidean space).

University of Washington: Martin Tompa; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352093 A01 & A02); $41,000; 12 mos.

This research involves the study of the computational complexity of various algorithms.

The following topics are being explored:

1. Probabilistic algorithms,
2. The parallel complexity of algebraic & number-theoretic problems,
3. Complexity issues in cryptosystem security,
4. Properties of 2-person pebbling, a combinatorial model of parallel computation, and
5. Algebraic simplification.

University of Washington; John Zahorjan; *Presidential Young Investigator Award (Computer Research)*; (DCR-8352098 A01); $57,500; 12 mos.

The purpose of this work is to investigate the feasibility of, and mechanisms and policies to support load sharing. The first step involves evaluating the conflicting influences of response time improvement achieved by balancing the load and response time degradation caused by task migration overhead. Because the questions to be answered are broad and the details poorly defined, this problem is best addressed by modeling. The goal here is to determine system characteristics (e.g., workload profile, communication channel speed, processor speeds) necessary for load balancing to be a reasonable approach. Once this has been done, implementation of specific mechanisms and policies will be used to find more detailed answers to the more detailed questions that can be posed. The goal here is to evaluate the relative performance of techniques for deciding when to move tasks, which tasks to move, and where to move them.

Univ. of Wisconsin - Madison; Deborah A. Joseph; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451387); $41,000; 12 mos.

Professor Joseph is investigating problems in the areas of abstract complexity theory and of design and analysis of algorithms. In the area of abstract complexity theory, these problems include proof techniques for the separation of complexity classes that address the question of whether certain structural properties of complexity classes may be independent of formal proof systems (such as Peano Arithmetic). This work continues earlier research which introduced techniques from the theory of nonstandard models which analyze the strength of the axioms necessary to resolve questions of program behavior and complexity and systematicallly investigates proof techniques (diagonalizations, priority arguments, and forcing arguments) which may be useful in resolving structural separation techniques.

In the area of design and analysis of algorithms, Professor Joseph is continuing to develop efficient geometric algorithms and motion planning algorithms for robotics. These algorithms compute paths of motion for objects (robot arms, manipulators, etc.) moving in a confined workspace.

University of Wisconsin - Madison; Udi Manber; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451397); $40,997; 12 mos.

Professor Manber is investigating problems in the area of design and implementation of distributed algorithms for networks of computers. Because it is expected that within a few years many professsionals and researchers will have access to powerful workstations that will mainly be used as personal computers, it is highly desirable to be able to distribute computation intensive jobs to several machines. The main problem with doing this is the difficulty in writing distributed programs. Professor Manber is investigating the following related problems:

1. Development of efficient distributed algorithms for a variety of problems,
2. Design of easy-to-use packages for implementing specific types of distributed algorithms,
3. Relations among several models of computation for multicomputers, and
4. Fault tolerance in an environment of cooperating workstations.

University of Wisconsin-Madison; Mary Vernon; *Presidential Young Investigator Award (Computer Research)*; (DCR-8451405); $25,000; 12 mos.

The goal of this research is to apply and extend modeling techniques for analyzing the performance of distributed systems. Distributed systems which are based on multicomputers (i.e., machines that are interconnected by a local area network), and systems which are based on multiprocessors (i.e., machines that share memory) are being explored. One aspect of this work involves applying recent work on graph models and timed Petri nets to the analysis of system performance. Another aspect of this work involves applying queueing network models to these problems. Performance problems of interest include evaluation of distributed algorithms, operating system algorithms for resource sharing, and advanced multiprocessor architectures for high-speed numeric and symbolic computation.

BLANK PAGE

Small Business Innovation Research

Technology Development of California, Inc.; Harry G. Heard; *A 64-Bit Mini-Supercomputer*; (DCR-8460153); $40,000; 6 mos.

Two attributes of the current supercomputer milieu reveal an important national need: (A) the known distribution of supercomputer resources in America at other than university centers of excellence, and, (B) the results of recent statistical analyses that identify important (typical) patterns of supercomputer usage by researchers. It is claimed that these data indicate that greater national progress could be made in science and technology today through the ready availability of a low cost mini-supercomputer rather than another Fifth Generation Supercomputer.

A modest-performance, stand-alone, 64-bit, vector-processing mini-supercomputer system, with a compatible register architecture, and a compatible instruction set, that is microcoded to execute the identical object code of the mainline supercomputer, could:

1. Substantially improve researcher software development productivity,
2. Increase the effectiveness of existing supercomputer facilities by releasing supercomputer resources to work more exclusively on compute intensive problems,
3. Provide more competitive resources for simulation and numerical analysis algorithm development, and
4. Improve the quality of training received by research scientists who could be trained, rather than retrained, to use scientific vector processors in pursuit of the computational experiments in advanced research.

Signatron, Inc.; Mark S. Wallace; *Consensus Theory and Expert Systems*; (DCR-8460485); $39,716; 6 mos.

By refining and deploying concepts and approaches drawn from the foundations of probability this project is developing additional methods for designing the knowledge base and inference mechanism of an expert system and evaluating existing methodologies for assigning weights to facts and rules obtained from experts. The concept of interval-valued probability is appropriate for modeling the imprecision and indeterminacy in the beliefs of an individual/expert and should allow us to aggregate the beliefs of several experts into a consensus, while preserving information as to the extent of disagreement. This represents a substantial step forward in realistic modeling and utilization of experts. The evaluation of existing methodologies (e.g., Bayesian, fuzzy sets, belief functions) has both an experimental component based upon a testbed expert system and an analytical framework patterned after similar analyses conducted in the foundations of probability.

William V. Stoecker; William V. Stoecker; *Skin Cancer Recognition by Computer Vision*; (DCR-8460921); $40,000; 6 mos.

A feasibility study is being undertaken for a computer vision device to help general practitioners in diagnosing skin cancers. The overall goal is to combine a currently available medical expert system computer program with a computer vision front end to create a useful diagnostic tool for skin cancers.

The research is developing model algorithms for a knowledge based feature detection system for 35 mm photographic slides of skin tumors. The research uses a VAX 11/780 computer and a Gould DeAnza IP-8400 image processing system. The hope is to detect the tumor features with highest tumor predictivity and to outline approaches to detect the remaining tumor features.

C. Abaci, Inc.; Edward L. Battiste; *Scientific Software for Micro-Computer Users: Quality Control*; (DCR-8313661 A01); $138,227; 18 mos.

The use of micro-computers is migrating from specific to general purpose. Hardware and software technology are evolving rapidly. No significant software effort is taking place so as to achieve the goals of providing robust software kernels and problem solving environments for scientific users and applying pressure toward the acceptance of standards for languages and software which will allow positive evolution of usage of micro-computers. The advances in software of the last decade are not transferring to micro-computer users.

The Phase I study investigated the feasibility of producing scientific software adequate for use by the (new) micro-computing audience. Problems were isolated which increase difficulty in, and cost of, production. Many will vanish with technology evolution. In Phase II the major problem encountered, quality control in a multi-computer, multi-language environment set with limited external storage facilities, is being

investigated. Embedded in that problem is a fundamental arithmetic problem. Most computer releases present incorrect arithmetic and elementary function facilities, to the extent that scientific use may be dangerous. Techniques which isolate such errors are being studied. The quality control effort utilizes a host-target computer concept in a local network. This can lead to non-local network implementations which will serve the dispersed using audience with robust scientific software.

Speech Response Systems, Inc.; Ronald A. Cole: *Research on Feature-Based Systems for Speech Recognition*; (DCR-8460412); $40,000; 6 mos.

The goal of this research is to determine the feasibility of developing a speaker-independent speech recognition system that can be implemented on currently available hardware components. The system uses a feature-based approach to speech recognition, in which the perceptually relevant features of speech are extracted from the signal and used to identify speaker-independent patterns for the consonant and vowel in an utterance. Previous research demonstrated that the feature-based approach is more powerful than the template matching approach used in all commercially available systems. To determine the feasibility of a real time system, it is necessary to develop a prototype system and then examine the computational requirements of the algorithms in light of currently available hardware components.

Index of Principal Investigators

Name — Page(s)

A

Adleman, Leonard ... 60
Agrawala, Ashok K. ... 78
Allen, James F. ... 31, 88
Amarel, Saul ... 30
Andrews, Gregory R. ... 47
Andrews, Peter B. ... 25
Angluin, Dana C. ... 24
Appelt, Douglas ... 21
Armstrong, James ... 12
Arnow, David M. ... 75
Ashcroft, Edward ... 10
Asmus, John F. ... 56
Atallah, Mikhail ... 86
Austing, Richard H. ... 57
Avizienis, Algirdas ... 13, 77

B

Babai, Laszlo ... 56
Babuska, Ivo ... 2
Bach, Eric ... 69
Baer, Jean-Loup ... 12
Bajcsy, Ruzena ... 20, 23, 81
Ballard, Dana H. ... 20
Balzer, Robert M. ... 36, 45
Banerji, Ranan B. ... 26
Barnhill, Robert E. ... 15, 34
Basili, Victor R. ... 78
Battiste, Edward L. ... 91
Ben-Bassat, Moshe ... 28
Berman, Piotr ... 65
Bernstein, Arthur J. ... 42
Bernstein, Herbert T. ... 19
Bhanu, Bir ... 21
Birman, Kenneth ... 35
Bitton, Dina ... 43
Bixby, Robert E. ... 3
Blass, Andreas R. ... 71
Bledsoe, Woodrow W. ... 26
Blum, Edward K. ... 39
Blum, Lenore ... 70
Blum, Manuel ... 51, 59
Book, Ronald V. ... 73
Boyer, Robert S. ... 26
Breuer, Melvin ... 9
Brown, Frank M. ... 49
Brown, Richard ... 54
Browne, James C. ... 82
Bruce, Kim B. ... 71
Buchanan, Bruce B. ... 27
Bunch, James R. ... 1
Buneman, Peter ... 81
Burks, Arthur W. ... 11, 29
Byrd, Richard H. ... 1

C

Carlson, David A. ... 75
Cartwright, Robert S. ... 82
Chang, Shi-Kuo ... 28
Cheriton, David R. ... 85
Clarke, Edmund M. ... 41
Clarke, Lori A. ... 34, 36
Cohen, Jacques ... 24, 38, 75
Cole, Richard ... 68
Cole, Ronald A. ... 92
Collins, George E. ... 7
Comer, Douglas E. ... 53, 78
Compton, Kevin J. ... 71
Constable, Robert L. ... 30, 37, 47, 80
Conte, Samuel ... 33
Conway, Richard W. ... 80
Cook, Robert ... 76
Cooper, Robert S. ... 10
Csuri, Charles A. ... 15
Cybenko, George ... 75

D

Dale, Alfred G. ... 82
Daley, Robert P. ... 73
Dasgupta, Subrata ... 46
Davida, George I. ... 51
Davis, Ernest ... 31
DeJong, Gerald ... 75
Demers, Alan ... 30, 80
Demmel, James ... 2
Dennis, John E. ... 82
Despain, Alvin M. ... 10
Dhall, Sudarshan K ... 57
Director, Stephen W. ... 9
Dobkin, David P. ... 14
Doeppner, Thomas W. ... 81
Dongarra, Jack ... 5
Dubinsky, Edward ... 57
Dunsmore, Hubert ... 33

E

Ehrenfeucht, Andrzej 73
Ehrlich, R. W. 34
Estrin, Gerald 77
Evens, Martha W. 22

F

Faustini, Antony 10
Feigenbaum, Edward A. 27
Feldman, Jerome A. 80
Fich, Faith E. 65
Fischer, Charles N. 48
Fischer, Michael J. 47
Flaherty, Joseph E. 76
Franklin, Mark A. 9
Frederickson, Greg N. 61
Fredman, Michael 60
Freeman, Peter A. 35
Friedman, Daniel P. 39
Furst, Merrick L. 68, 88

G

Gabow, Harold N. 66
Gafni, Eliezer 86
Galil, Zvi 63
Gallager, Robert G. 53
Galotti, Kathleen 54
Gautschi, Walter 1
Gaylord, Arthur S. 18, 33, 54
Gear, Charles W. 78
Giacalone, Alessandro 76
Gilbert, John 88
Ginsburg, Seymour 66
Goldwasser, Shafrira 62
Golub, Gene H. 1
Good, Donald I. 82
Gottlieb, Allan J. 11
Grant, John 25
Greenberg, Donald P. 15
Grenander, Ulf 26
Gries, David J. 47, 75, 80
Grishman, Ralph E. 22
Griswold, Ralph E. 37, 38
Grosse, Eric 5
Grunbaum, Alberto 17
Gurevich, Yuri 71
Guttag, John V. 40, 70

H

Hanson, David R. 77
Harrison, Michael 60
Hartigan, John A. 5
Hartmanis, Juris 63
Hartson, H. R. 34
Hayes, John P. 13
Haynes, Christopher T. 39
Heard, Harry G. 91
Heller, Jack 80
Hellman, Martin E. 51
Henderson, Thomas C. 21
Hennessy, John L. 85
Henry, Robert 37
Henschen, Lawrence J. 24
Hilfinger, Paul N. 85
Hirshman, Lynette 23
Hofstadter, Douglas R. 29
Hollaar, Lee A. 12, 82
Hopcroft, John E. 30
Hsiang, Jieh 76
Huang, Thomas S. 17
Hudak, Paul R. 28, 86
Hughes, Herman D. 54
Humblet, Pierre A. 53
Hummel, Robert 19
Hunt, Harry B. 74
Hutchinson, Joan P. 70

I

Ibarra, Oscar H. 62
Imielinski, Tomasz 25
Irwin, Mary Jane 76

J

Jardetzky, Oleg 27
John, David J. 55
Johnson, James L. 55
Johnson, Steven D. 46
Joseph, Deborah A. 73, 89
Joshi, Aravind K. 23, 81
Jump, J. Robert 82

K

Kafura, Dennis G. 34
Kak, Avinash 17
Kaltofen, Erich 7
Kamimura, Tsutomu 39
Kanade, Takeo 20
Kanal, Laveen N. 29
Kannan, Ravindran 6, 64
Kant, Elaine 25
Karp, Richard M. 59, 60
Karplus, Kevin 9
Karplus, Walter 77
Kedem, Zvi M. 55
Kender, John R. 87
Kennedy, Kenneth W. 37, 82
Kessler, Robert R. 38
Kincaid, David R. 6
Kirby, L. A. S. 71
Kirchner, Roger 54
Klee, Victor L. 65
Kleinrock, Leonard 77

Klema, Virginia C. 5
Knuth, Donald E. 45
Ko, Ker-I 65
Komlos, Janos 66
Korb, J. Tim 53, 78
Korth, Henry F. 44
Kosaraju, S. Rao 62
Kreiss, Heinz-Otto 1
Kubitz, W. 78
Kuck, David J. 11, 39
Kuipers, Benjamin J. 31

L

Ladner, Richard E. 69
Lakshmivarahan, S. 57
Lam, Simon 55
Landau, Susan 70
Landweber, Lawrence H. 83
Lawler, Eugene L. 60
Lawrie, Duncan H. 11, 39
Lazowska, Edward D. 82
Lee, Der-Tsai 61
Lee, Insup 23
Leighton, Frank T. 87
Leone, James A. 55
Lesser, Victor R. 14, 29, 53, 78
Leveson, Nancy G. 34
Levin, Leonid A. 70
Lewis, Harry R. 70
Lindstrom, Gary 41, 76
Lipton, Richard 79
Liu, C. L. 61, 78
Lowe, David 19
Lueker, George S. 60
Luks, Eugene 64
Lynch, James F. 72
Lynch, Nancy 67

M

Ma, Yuen-Wah Eva 88
Maass, Wolfgang 61
Manber, Udi 89
Mangasarian, Olvi 3
Manna, Zohar 45, 66
Marek, Wiktor 43, 56
Martel, Charles U. 60
Matsen, Frederick A. 6
Matula, David W. 12
Mayr, Ernst W. 85
McAloon, Kenneth 68
McCarthy, John 23
McCluskey, Edward J. 13
McCracken, Nancy J. 41
McDermott, Drew V. 28
McKendry, Martin S. 48
Meyer, Carl 3
Meyer, Robert R. 4
Micali, Silvio 56, 62
Michalski, R. S. 24
Miller, Gary L. 60, 62
Minker, Jack 25, 78
Mitchell, Tom M. 87
Moore, J. Strother 26
Morris, Robert 57
Moser, Louise E. 45
Mueller, Robert 38
Mulvey, John M. 6
Murray, Walter 4
Musser, David R. 40, 72
Myers, Eugene W. 38

N

Nau, Dana S. 87
Nau, Richard 54
Ness, Linda 54
Nevins, Arthur J. 24
Newell, Allen 25
Nicolau, Alexandru 40
Nirenberg, Robert 40
Notkin, David S. 37

O

O'Dunlaing, Colm 31
O'Neil, Elizabeth J. 54
O'Rourke, Joseph 87
Ordman, Edward 68
Osterweil, Leon 36, 77
Overton, Michael 3
Owens, Robert 76
Ozsoyoglu, Gultekin 44
Ozsoyoglu, Z. Meral 44

P

Pan, Victor Y. 64
Papadimitriou, Christos H. 65
Parikh, Rohit J. 72
Parker, Alice 9
Patterson, Wayne 53
Pearl, Judea 27, 28
Perlis, Alan J. 77
Perros, Harry 14
Peters, Stanley 56
Plemmons, Robert 3
Pomerance, Carl 61
Popek, Gerald 77
Poplawski, David A. 54
Pratt, Terrence 37

R

Rabin, Michael O. 66
Ramamritham, Krithivasan 78
Reid, Brian K. 85
Reif, John 67

Rennels, David A. 13
Reynolds, John C. 41
Rheinboldt, Werner C. 3, 6
Rice, John 75
Richards, Whitman 18
Richardson, Debra J. 34
Riddle, William E. 36, 44
Riesenfeld, Richard F. 82
Riseman, Edward M. 18, 33, 54, 78
Rivest, Ronald L. 62
Robertson, Edward 75
Robinson, Stephen 4
Romney, Leonard C. 83
Rose, Donald J. 79
Rosen, J. B. 2
Rosenfeld, Azriel 18, 78
Rowe, Lawrence A. 35
Rudd, Walter G. 11
Russell, Robert 2
Ruzzo, Walter 69

S

Sahni, Sartaj K. 63, 79
Sameh, Ahmed 5
Samet, Hanan 18
Savage, John E. 68, 81
Schank, Roger C. 77
Schlicting, Richard D. 47
Schmidt, Charles 30
Schnabel Robert 77
Schneider, Fred B. 47
Schnitger, George 65
Schultz, Martin H. 77
Schwartz, Jacob T. 11, 19, 80
Sciore, Edward 44
Sedelow, Sally Y. 21
Sedelow, Walter A. 21
Sedgewick, Robert 46, 47, 65, 67
Selman, Alan L. 74
Seth, Sharad C. 63
Sethi, Ishwar K. 18
Shafer, Steven A. 20
Shasha, Dennis 42
Shen, Vincent 33
Siewiorek, Daniel P. 9
Silberschatz, Abraham 44, 82
Sipser, Michael 62
Slagle, James R. 30
Slomam, Theodore A. 86
Slater, Peter J. 59
Sloan, Jr., Kenneth R. 15
Smale, Stephen 59
Smith, Alan J. 42
Smith, Jeffrey W. 61
Smolka, Scott A. 42, 76
Snodgrass, Richard T. 40
Snyder, Lawrence 48
Soare, Robert 56
Soloway, Elliot M. 33
Solworth, Jon 11
Spirakis, Paul 43
Srini, Vason P. 10
Srivas, Mandayam 76
Stanat, Donald F. 79
Stearns, Richard 74
Stemple, David W. 43
Stewart, Gilbert W. 78
Stewart, William 14, 33
Stoecker, William V. 91
Storer, James A. 53
Stout, Quentin F. 67
Switzer, Paul 17
Szilagyi, Mike 4

T

Tang, Adrian 39
Tanimoto, Steven L. 21
Taylor, Richard 35, 86
Thomas, Donald 9
Thompson, William B. 19
Tichy, Walter 53, 78
Tompa, Martin 7, 69, 89
Traub, Joseph F. 72
Treu, Siegfried 52
Tripp, Edwin 15
Tucker, Alan 64
Tucker, Allen B. 22
Tuler, Randy 61
Turchin, Valentin 40
Turner, Jonathan S. 67

U

Uhr, Leonard 13

V

Valiant, Leslie G. 25
van Dam, Andries 81
Vazirani, Vijay 63
Vernon, Mary 49, 89
Vishkin, Uzi 14
Vitter, Jeffrey S. 65, 88

W

Wahba, Grace 4
Waldinger, Richard J. 45
Walker, Terry M. 75
Wallace, Mark S. 91
Walter, Gilbert G. 51
Wang, Kuo K. 30
Wang, Paul S. 7
Wann, Donald F. 9
Warren, David S. 44
Wegner, Peter 81

Weihl, William E. 48
Wells, Mark B. 40
Welsch, Roy E. 5
Weyuker, Elaine 35
Widlund, Olof 3, 76, 80
Wile, David 36, 45
Wileden, Jack C. 36
Willard, Dan E. 68
Winklmann, Karl 61
Wise, David S. 46

Y

Yanofsky, Charles 27
Yao, Andrew C. 59
Yap, Chee-Keng 31, 68
Ylvisaker, N. Donald 5
Young, David M. 6
Young, Paul R. 58
Young, Tzay Y. 17
Yun, David Y. 7

Z

Zadeh, Lotfi A. 27
Zahorjan, John 89
Zeigler, Bernard P. 52
Zeil, Steven J. 34
Zorat, Alessandro 76
Zucker, Jeffery I. 41